

INTERNATIONAL
STANDARD

ISO/
IEC/IEEE
8802-3

Third edition
2021-02

AMENDMENT 1
2021-12

**Telecommunications and exchange
between information technology
systems — Requirements for local and
metropolitan area networks —**

Part 3:
Standard for Ethernet

**AMENDMENT 1: Physical layer
specifications and management
parameters for 2.5 Gb/s and 5 Gb/s
operation over backplane**

*Télécommunications et échange entre systèmes informatiques —
Exigences pour les réseaux locaux et métropolitains —*

Partie 3: Norme pour Ethernet

*AMENDEMENT 1: Spécifications de la couche physique et
paramètres de gestion pour fonctionnement à 2,5 Gb/s et 5 Gb/s sur
fond de panier*



Reference number
ISO/IEC/IEEE 8802-3:2021/Amd.1:2021(E)

© IEEE 2019



COPYRIGHT PROTECTED DOCUMENT

© IEEE 2019

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

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

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

Published in Switzerland

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.

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

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

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

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

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

ISO/IEC/IEEE 8802-3:2021/Amd.1 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE Std 802.3cb-2018) and drafted in accordance with its editorial rules. It was adopted, under the “fast-track procedure” defined in the Partner Standards Development Organization cooperation agreement between ISO and IEEE, by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*.

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

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

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

IEEE Std 802.3cb™-2018
(Amendment to IEEE Std 802.3™-2018)

IEEE Standard for Ethernet

Amendment 1: Physical Layer Specifications and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation over Backplane

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 27 September 2018

IEEE-SA Standards Board

STANDARDSISO.COM: Click to view the full PDF of ISO/IEC JTC1/SC29/WG6/8802-3:2021/Amd 1:2021

Abstract: Ethernet Media Access Control (MAC) parameters, Physical Layer specifications, and management objects for the serial transfer of Ethernet format frames at 2.5 Gb/s and 5 Gb/s over electrical backplanes are defined in this amendment to IEEE Std 802.3-2018.

Keywords: 2.5 Gigabit Ethernet, 5 Gigabit Ethernet, 2.5GBASE-KX, 2.5GBASE-X, 5GBASE-KR, 5GBASE-R, amendment, AN, Auto-Negotiation, Backplane Ethernet, BASE-R, BASE-X, EEE, Energy Efficient Ethernet, Ethernet, IEEE 802.3™, IEEE 802.3cb™

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

Copyright © 2019 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 4 January 2019. Printed in the United States of America.

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

PDF: ISBN 978-1-5044-5278-6 STD23391
Print: ISBN 978-1-5044-5279-3 STDPD23391

IEEE prohibits discrimination, harassment and bullying.

For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

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

Important Notices and Disclaimers Concerning IEEE Standards Documents

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

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

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

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

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

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

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

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

Translations

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

Official statements

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

Comments on standards

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

Comments on standards should be submitted to the following address:

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

Laws and regulations

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

Copyrights

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

Photocopies

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

Updating of IEEE Standards documents

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

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

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

Errata

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

Patents

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

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

Participants

The following individuals were officers and members of the IEEE 802.3 Working Group at the beginning of the IEEE P802.3cb Working Group ballot.

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

Yong Kim, *IEEE P802.3cb Task Force Chair, Phase 1**
Daniel F. Smith, *IEEE P802.3cb Task Force Chair, Phase 2*
William Lo, *IEEE P802.3cb Task Force Editor-in-Chief, Phase 1*
Daniel F. Smith, *IEEE P802.3cb Task Force Editor-in-Chief, Phase 2*
Jim Hatfield, *IEEE P802.3cb Task Force Editor, Phase 2**

Justin Abbott	John Dillard	Upen Kareti
David Abramson	Daniel Dillow	Keisuke Kawahara
Shadi Abughazaleh	Thuyen Dinh	Yasuaki Kawatsu
Mohammad Ahmed	Curtis Donahue	Michael Kelsen
Eric Baden	Dan Dove	Scott Kipp
Amrik Bains	Mike Dudek	Michael Klempa
Thananya Baldwin	David Dwelley	Curtis Knittle
Denis Beaudoin	Frank Effenberger	Shigeru Kobayashi
Christian Beia	Hesham Elbakoury	Daniel Koehler
Michael Bennett	David Estes	Paul Kolesar
Vipul Bhatt	John Ewen	Tom Kolze
William Bliss	Ramin Farjad	Glen Kramer
Brad Booth	Shahar Feldman	Hans Lackner
Martin Bouda	James Fife	Jeffrey Lapak
Ralf-Peter Braun	Alan Flatman	Mark Laubach
Theodore Brillhart	Matthias Fritsche	Han Hyub Lee
Paul Brooks	Richard Frosch	David Lewis
Alan Brown	Andrew Gardner	Jon Lewis
Matthew Brown	Claude Gauthier	Mike Peng Li
Chris Bullock	Ali Ghiasi	Jane Lim
Jairo Bustos Heredia	Joel Goergen	Dekun Liu
Adrian Butter	Volker Goetzfried	Hai-Feng Liu
Francesco Caggioni	Zhigang Gong	Miklos Lukacs
Anthony Calbone	Steven Gorshe	Kent Lusted
Clark Carty	Robert Grow	Jeffery Maki
Craig Chabot	Mark Gustlin	David Malicoat
Geoffrey Chacon Simon	Marek Hajduczenia	Yonatan Malkiman
Mandeep Chadha	Takehiro Hayashi	Arthur Marris
David Chalupsky	Yasuo Hidaka	Takeo Masuda
Jacky Chang	Rita Horner	Erdem Matoglu
Xin Chang	Bernd Hormmeyer	Naoki Matsuda
Ahmad Chini	Victor Hou	Mick McCarthy
Keng Hua Chuang	Yasuhiro Hyakutake	Brett McClellan
Christopher R. Cole	Hideki Isono	Thomas McDermott
John D'Ambrosia	Tom Issenhuth	John McDonough
Yair Darshan	Kenneth Jackson	Larry McMillan
Piers Dawe	Andrew Jimenez	Richard Mei
Fred Dawson	Chad Jones	Richard Mellitz
Wael Diab	Peter Jones	Bryan Moffitt
Eric DiBiaso	Manabu Kagami	Ardeshir Mohammadian
		Paul Mooney

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

Dale Murray	Sam Sambasivan	Pirooz Toyserkani
Henry Muysshondt	Edward Sayre	Albert Tretter
James Nadolny	Dieter Schicketanz	Stephen Trowbridge
Edward Nakamoto	Fred Schindler	Yoshihiro Tsukamoto
Gary Nicholl	Hossein Sedarat	Ed Ulrichs
Kevin Noll	Naoshi Serizawa	Alexander Umnov
Mark Nowell	Masood Shariff	Sterling A. Vaden
David Ofelt	Ramin Shirani	Stefano Valle
Thomas Palkert	Tom Skaar	Paul Vanderlaan
Hui Pan	Jeff Slavick	Robert Wagner
Sesha Panguluri	Scott Sommers	Dylan Walker
Vasu Parthasarathy	Yoshiaki Sone	Haifei Wang
Petar Pepeljugin	Tom Souvignier	Roy Wang
Gerald Pepper	Peter Stassar	Tongtong Wang
Ruben Perez De Aranda Alonso	Heath Stewart	Xinyuan Wang
Michael Peters	Robert Stone	Matthias Wendt
Phong Pham	David Stover	Oded Wertheim
Jean Picard	Junqing Sun	Natalie Wienckowski
William Powell	Ken-Ichi Suzuki	Ludwig Winkel
Rick Rabinovich	Steve Swanson	Peter Wu
Adee Ran	Andre Szczepanek	Dayin Xu
Alon Regev	William Szeto	Yu Xu
Duane Remein	Bharat Tailor	Jun Yi
Victor Renteria	Takayuki Tajima	Lennart Yseboodt
Christopher Roth	Satoshi Takahashi	Hayato Yuki
Salvatore Rotolo	Kohichi Tamura	Andrew Zambell
Toshiaki Sakai	Brian Teipen	Yan Zhuang
Jorge Salinger	Geoffrey Thompson	George Zimmerman

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

Shadi Abughazaleh	Marco Hernandez	Thomas Palkert
Thomas Alexander	David Hess	Bansi Patel
Peter Anslow	Werner Hoelzl	Arumugam Paventhan
Butch Anton	Rita Horner	Michael Peters
Amrik Bains	C. Huntley	Rick Pimpinella
Ralf-Peter Braun	Noriyuki Ikeuchi	Adee Ran
Nancy Bravin	Atsushi Ito	Alon Regev
Theodore Brillhart	Raj Jain	Maximilian Riegel
Matthew Brown	SangKwon Jeong	Robert Robinson
Demetrio Bucaneg	Piotr Karocki	Benjamin Rolfe
Jairo Bustos Heredia	Stuart Kerry	Daniel F. Smith
William Byrd	Yong Kim	Thomas Starai
Steven B. Carlson	Mark Laubach	Walter Struppler
Juan Carreon	David J. Law	Mitsutoshi Sugawara
Mandeep Chadha	Jon Lewis	Junqing Sun
Keith Chow	Mike Peng Li	Geoffrey Thompson
Keng Hua Chuang	Elvis Maculuba	Michael Thompson
Charles Cook	Valerie Maguire	Steven Tilden
Rodney Cummings	Jeffery Maki	David Tremblay
Sourav Dutta	Arthur Marris	Mark-Rene Uchida
Michael Fischer	Mick McCarthy	Alexander Umnov
Matthias Fritsche	Brett McClellan	Paul Vanderlaan
Yukihiro Fujimoto	Thomas McDermott	George Vlantis
Zhigang Gong	Larry McMillan	Khurram Waheed
James Graba	Richard Mellitz	Lisa Ward
Eric W. Gray	Charles Moorwood	Karl Weber
Randall Groves	Michael Newman	Andreas Wolf
Robert Grow	Nick S. A. Nikjoo	Peter Wu
Michael Gundlach	Paul Nikolich	Jun Xu
Marek Hajduczenia	Satoshi Obara	Oren Yuen
Adam Healey		Zhen Zhou

When the IEEE-SA Standards Board approved this amendment on 27 September 2018, it had the following membership:

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

Ted Burse
Guido R. Hiertz
Christel Hunter
Joseph L. Koepfinger*
Thomas Koshy
Hung Ling
Dong Liu

Xiaohui Liu
Kevin Lu
Daleep Mohla
Andrew Myles
Paul Nikolich
Ronald C. Petersen
Annette D. Reilly

Robby Robson
Dorothy Stanley
Mehmet Ulema
Phil Wemblom
Philip Winston
Howard Wolfman
Jingyi Zhou

*Member Emeritus

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

Introduction

This introduction is not part of IEEE Std 802.3cb-2018, IEEE Standard for Ethernet—Amendment 1: Physical Layer Specifications and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation over Backplane.

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 802.3-2018 and are not maintained as separate documents.

At the date of publication for IEEE Std 802.3cb-2018, IEEE Std 802.3 was composed of the following documents:

IEEE Std 802.3-2018

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.

Section Seven—Includes Clause 96 through Clause 115 and Annex 97A through Annex 115A. Clause 96 through Clause 98, Clause 104, and associated annexes, specify Physical Layers and optional features for 100 Mb/s and 1000 Mb/s operation over a single twisted pair. Clause 100 through Clause 103, as well as associated annexes, specify Physical Layers for the operation of the EPON protocol over coaxial distribution networks. Clause 105 through Clause 114 and associated annexes include general information on 25 Gb/s operation as well as 25 Gb/s Physical Layer specifications. Clause 99 specifies a MAC merge sublayer for the interspersing of express traffic. Clause 115 and its associated annex specify a Physical Layer for 1000 Mb/s operation over plastic optical fiber.

Section Eight—Includes Clause 116 through Clause 126 and Annex 119A through Annex 120E. Clause 116 through Clause 124 and associated annexes include general information on 200 Gb/s and 400 Gb/s operation as well the 200 Gb/s and 400 Gb/s Physical Layer specifications. Clause 125 and Clause 126 include general information on 2.5 Gb/s and 5 Gb/s operation as well as 2.5 Gb/s and 5 Gb/s Physical Layer specifications.

IEEE Std 802.3cb-2018

Amendment 1—This amendment includes changes to IEEE Std 802.3-2018 and its amendments, and adds Clause 127 through Clause 130, Annex 127A, Annex 128A, Annex 128B, and Annex 130A. This amendment adds new Physical Layers for operation at 2.5 Gb/s and 5 Gb/s over electrical backplanes.

Two companion documents exist, IEEE Std 802.3.1 and IEEE Std 802.3.2. IEEE Std 802.3.1 describes Ethernet management information base (MIB) modules for use with the Simple Network Management Protocol (SNMP). IEEE Std 802.3.2 describes YANG data models for Ethernet. IEEE Std 802.3.1 and IEEE Std 802.3.2 are updated to add management capability for enhancements to IEEE Std 802.3 after approval of those enhancements.

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

Contents

1. Introduction.....	23
1.3 Normative references	23
1.4 Definitions	23
1.5 Abbreviations.....	24
30. Management.....	25
30.3 Layer management for DTEs.....	25
30.3.2 PHY device managed object class	25
30.3.2.1 PHY device attributes	25
30.3.2.1.2 aPhyType	25
30.3.2.1.3 aPhyTypeList	25
30.3.2.1.5 aSymbolErrorDuringCarrier	25
30.5 Layer management for medium attachment units (MAUs).....	26
30.5.1 MAU managed object class	26
30.5.1.1 MAU attributes	26
30.5.1.1.2 aMAUType	26
30.6 Management for link Auto-Negotiation	26
30.6.1 Auto-Negotiation managed object class	26
30.6.1.1 Auto-Negotiation attributes	26
30.6.1.1.5 aAutoNegLocalTechnologyAbility	26
45. Management Data Input/Output (MDIO) Interface.....	27
45.2 MDIO Interface Registers.....	27
45.2.1 PMA/PMD registers	27
45.2.1.1 PMA/PMD control 1 register (Register 1.0).....	27
45.2.1.1.5 PMA local loopback (1.0.0).....	27
45.2.1.6 PMA/PMD control 2 register (Register 1.7).....	27
45.2.1.7 PMA/PMD status 2 register (Register 1.8).....	28
45.2.1.7.4 Transmit fault (1.8.11).....	28
45.2.1.7.5 Receive fault (1.8.10).....	28
45.2.1.8 PMD transmit disable register (Register 1.9)	28
45.2.1.18 2.5G/5G PMA/PMD extended ability register (Register 1.21)	29
45.2.1.18.a 5GBASE-KR ability (1.21.3).....	29
45.2.1.18.b 2.5GBASE-KX ability (1.21.2)	29
45.2.1.89 BASE-R PMD control register (Register 1.150)	29
45.2.1.89.3 Transmitter equalizer disable (1.150.2)	30
45.2.1.97 1000BASE-KX/2.5GBASE-KX control register (Register 1.160)	30
45.2.1.97.1 PMD transmit disable (1.160.0).....	30
45.2.1.98 1000BASE-KX/2.5GBASE-KX status register (Register 1.161).....	30
45.2.1.98.6 1000BASE-KX/2.5GBASE-KX signal detect (1.161.0).....	31
45.2.1.140 Test-pattern ability (Register 1.1500).....	31
45.2.3 PCS registers.....	31
45.2.3.1 PCS control 1 register (Register 3.0).....	31
45.2.3.1.2 Loopback (3.0.14).....	31
45.2.3.6 PCS control 2 register (Register 3.7).....	32
45.2.3.8 PCS status 3 register (Register 3.9).....	32
45.2.3.8.a 5GBASE-R capable (3.9.3).....	32
45.2.3.8.b 2.5GBASE-X capable (3.9.2).....	32

45.2.3.11	EEE control and capability 2 register (Register 3.21)	33
45.2.3.11.a	5GBASE-KR EEE supported (3.21.8)	33
45.2.3.11.b	2.5GBASE-KX EEE supported (3.21.7)	33
45.2.3.15	BASE-R and MultiGBASE-T PCS status 1 register (Register 3.32)	33
45.2.3.15.1	BASE-R and MultiGBASE-T receive link status (3.32.12)	33
45.2.3.15.4	BASE-R and MultiGBASE-T PCS high BER (3.32.1)	34
45.2.3.15.5	BASE-R and MultiGBASE-T PCS block lock (3.32.0)	34
45.2.3.16	BASE-R and MultiGBASE-T PCS status 2 register (Register 3.33)	34
45.2.3.16.1	Latched block lock (3.33.15)	34
45.2.3.16.2	Latched high BER (3.33.14)	34
45.2.3.16.3	BER(3.33.13:8)	34
45.2.3.16.4	Errored blocks (3.33.7:0)	35
45.2.3.17	5/10/25GBASE-R PCS test pattern seed A (Registers 3.34 through 3.37)	35
45.2.3.18	5/10/25GBASE-R PCS test pattern seed B (Registers 3.38 through 3.41)	35
45.2.7	Auto-Negotiation registers	35
45.2.7.2	AN status (Register 7.1)	35
45.2.7.2.1	Parallel detection fault (7.1.9)	35
45.2.7.12	Backplane Ethernet, BASE-R copper status (Register 7.48)	36
45.2.7.12.3	Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, 7.48.5, 7.48.6, 7.48.8, 7.48.9, 7.48.10, 7.48.11, 7.48.12, 7.48.13, 7.48.14, 7.48.15)	36
45.2.7.15	EEE advertisement 2 (Register 7.62)	36
45.2.7.15.a	5GBASE-KR EEE (7.62.3)	37
45.2.7.15.b	2.5GBASE-KX EEE (7.62.2)	37
45.2.7.16	EEE link partner ability 2 (Register 7.63)	37
45.5	Protocol implementation conformance statement (PICS) proforma for Clause 45, Management Data Input/Output (MDIO) interface	38
45.5.3	PICS proforma tables for the Management Data Input Output (MDIO) interface	38
45.5.3.2	PMA/PMD MMD options	38
45.5.3.3	PMA/PMD management functions	38
45.5.3.6	PCS options	38
46.	Reconciliation Sublayer (RS) and 10 Gigabit Media Independent Interface (XGMII)	39
46.3	XGMII functional specifications	39
46.3.3	Error and fault handling	39
46.3.3.3	Response to received invalid frame sequences	39
69.	Introduction to Ethernet operation over electrical backplanes	40
69.1	Overview	40
69.1.1	Scope	40
69.1.2	Relationship of Backplane Ethernet to the ISO OSI reference model	40
69.2	Summary of Backplane Ethernet Sublayers	41
69.2.3	Physical Layer signaling systems	41
69.3	Delay constraints	42
69.5	Protocol implementation conformance statement (PICS) proforma	42
70.	Physical Medium Dependent sublayer and baseband medium, type 1000BASE-KX	43
70.5	PMD MDIO function mapping	43
73.	Auto-Negotiation for backplane and copper cable assembly	44
73.2	Relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model	44

73.6	Link codeword encoding	44
73.6.4	Technology Ability Field.....	44
73.7	Receive function requirements	45
73.7.4	Arbitration function requirements	45
73.7.4.1	Parallel Detection function	45
73.7.6	Priority Resolution function.....	45
73.10	State diagrams and variable definitions	46
73.10.1	State diagram variables.....	46
73.10.2	State diagram timers	46
73.11	Protocol implementation conformance statement (PICS) proforma for Clause 73, Auto-Negotiation for backplane and copper cable assembly.....	47
73.11.4	PICS proforma tables for Auto-Negotiation for backplane and copper cable assembly....	47
73.11.4.4	Receive function requirements	47
78.	Energy-Efficient Ethernet (EEE).....	48
78.1	Overview.....	48
78.1.1	LPI Signaling.....	48
78.1.4	PHY types optionally supporting EEE	48
78.2	LPI mode timing parameters description.....	48
78.5	Communication link access latency.....	49
125.	Introduction to 2.5 Gb/s and 5 Gb/s networks	50
125.1	Overview.....	50
125.1.2	Relationship of 2.5 Gigabit and 5 Gigabit Ethernet to the ISO OSI reference model.....	50
125.1.3	Nomenclature.....	50
125.1.4	Physical Layer signaling systems	51
125.2	Summary of 2.5 Gigabit and 5 Gigabit Ethernet sublayers	52
125.2.2	Physical coding sublayer (PCS).....	52
125.2.3	Physical Medium Attachment sublayer (PMA).....	52
125.2.4	Auto-Negotiation, type BASE-T	52
125.2.4.1	Auto-Negotiation, type BASE-T	52
125.2.4.2	Auto-Negotiation, type Backplane	53
125.3	Delay Constraints.....	53
127.	Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X	54
127.1	Overview.....	54
127.1.1	Scope.....	54
127.1.2	Relationship of 2.5GBASE-X to other standards	54
127.1.3	Summary of 2.5GBASE-X sublayers	54
127.1.3.1	Physical Coding Sublayer (PCS).....	54
127.1.3.2	Physical Medium Attachment (PMA) sublayer.....	54
127.1.3.3	Physical Medium Attachment (PMA) service interface rates	55
127.1.4	Inter-sublayer interfaces	55
127.1.5	Functional block diagram	56
127.2	Physical Coding Sublayer (PCS).....	56
127.2.1	PCS Interface (XGMII)	56
127.2.2	Functions within the PCS	57
127.2.3	PCS used with 2.5GBASE-KX PMD	57
127.2.4	Use of code-groups	58
127.2.5	XGMII to 2.5GPPII mapping.....	58

127.2.5.1	2.5 Gb/s PCS Internal Interface (2.5GPII).....	58
127.2.5.2	Word Encode	59
127.2.5.3	Word-to-Octets	61
127.2.5.4	Octets-to-Word	61
127.2.5.5	Word Decode	62
127.2.6	8B/10B transmission code	62
127.2.6.1	Notation conventions	62
127.2.6.2	Transmission order	62
127.2.6.3	Generating code-groups and checking the validity of received code	64
127.2.6.4	Ordered sets	64
127.2.6.5	Comma considerations.....	65
127.2.6.6	Sequence (/Q/)	65
127.2.6.7	Data (/D/)	65
127.2.6.8	IDLE (/I/)	65
127.2.6.9	Low Power Idle (/LI/).....	65
127.2.6.10	Start_of_Packet delimiter (SPD)	65
127.2.6.11	End_of_Packet delimiter (EPD)	66
127.2.6.12	Error_Propagation (/V/).....	66
127.2.7	Detailed functions and state diagrams	66
127.2.7.1	State variables	66
127.2.7.1.1	Notation conventions	66
127.2.7.1.2	Constants.....	66
127.2.7.1.3	Variables	67
127.2.7.1.4	Functions.....	72
127.2.7.1.5	Counters	73
127.2.7.1.6	Messages.....	74
127.2.7.1.7	Timers	75
127.2.7.2	State diagrams.....	76
127.2.7.2.1	Word Encode and Word-to-Octets	76
127.2.7.2.2	Transmit.....	76
127.2.7.2.3	Synchronization.....	80
127.2.7.2.4	Receive.....	80
127.2.7.2.5	Octets-to-Word and Decode	80
127.2.7.2.6	LPI state diagram	86
127.2.7.2.7	LPI status and management	87
127.3	Physical Medium Attachment (PMA) sublayer.....	87
127.3.1	Service Interface	87
127.3.1.1	PMA_UNITDATA.request.....	87
127.3.1.1.1	Semantics of the service primitive.....	87
127.3.1.1.2	When generated	87
127.3.1.1.3	Effect of receipt	87
127.3.1.2	PMA_UNITDATA.indication	88
127.3.1.2.1	Semantics of the service primitive.....	88
127.3.1.2.2	When generated	88
127.3.1.2.3	Effect of receipt	88
127.3.2	Functions within the PMA	88
127.3.2.1	Data delay	88
127.3.2.2	PMA transmit function	88
127.3.2.3	PMA receive function.....	89
127.3.2.4	Code-group alignment	89
127.3.3	Loopback mode.....	89
127.3.3.1	Receiver considerations	89
127.3.3.2	Transmitter considerations.....	89
127.3.4	Test functions.....	89

127.3.4.1	PMA PRBS9 test pattern (optional).....	90
127.4	Compatibility considerations	90
127.5	Delay constraints.....	90
127.6	Environmental specifications.....	90
127.7	Protocol implementation conformance statement (PICS) proforma for Clause 127, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X.....	91
127.7.1	Introduction.....	91
127.7.2	Identification.....	91
127.7.2.1	Implementation identification.....	91
127.7.2.2	Protocol summary.....	91
127.7.3	Major capabilities/options.....	92
127.7.4	PICS proforma tables for the PCS and PMA sublayer, type 2.5GBASE-X.....	92
127.7.4.1	PCS.....	92
127.7.4.2	Code-group functions	93
127.7.4.3	EEE.....	93
127.7.4.4	PMA functions.....	94
127.7.4.5	Compatibility considerations	94
128.	Physical Medium Dependent sublayer and baseband medium, type 2.5GBASE-KX.....	95
128.1	Overview.....	95
128.2	Physical Medium Dependent (PMD) service interface	95
128.2.1	PMD_UNITDATA.request.....	96
128.2.1.1	Semantics of the service primitive.....	96
128.2.1.2	When generated	96
128.2.1.3	Effect of receipt	96
128.2.2	PMD_UNITDATA.indication	96
128.2.2.1	Semantics of the service primitive.....	96
128.2.2.2	When generated	96
128.2.2.3	Effect of receipt	96
128.2.3	PMD_SIGNAL.indication	96
128.2.3.1	Semantics of the service primitive.....	96
128.2.3.2	When generated	97
128.2.3.3	Effect of receipt	97
128.2.4	PMD_RXQUIET.request.....	97
128.2.4.1	Semantics of the service primitive.....	97
128.2.4.2	When generated	97
128.2.4.3	Effect of receipt	97
128.2.5	PMD_TXQUIET.request.....	97
128.2.5.1	Semantics of the service primitive.....	97
128.2.5.2	When generated	97
128.2.5.3	Effect of receipt	98
128.3	PCS requirements for Auto-Negotiation (AN) service interface.....	98
128.4	Delay constraints.....	98
128.5	PMD MDIO function mapping.....	98
128.6	PMD functional specifications.....	99
128.6.1	Link block diagram.....	99
128.6.2	PMD transmit function	100
128.6.3	PMD receive function.....	100
128.6.4	PMD signal detect function	100
128.6.5	PMD transmit disable function.....	100
128.6.6	Loopback mode.....	100
128.6.7	PMD fault function	101

128.6.8	PMD transmit fault function.....	101
128.6.9	PMD receive fault function.....	101
128.6.10	PMD LPI function	101
128.7	2.5GBASE-KX electrical characteristics.....	102
128.7.1	Transmitter characteristics.....	102
128.7.1.1	Test fixtures	102
128.7.1.2	Test fixture characteristics	102
128.7.1.3	Signaling speed	103
128.7.1.4	Output amplitude	103
128.7.1.5	Differential output return loss.....	104
128.7.1.6	Common-mode output return loss	104
128.7.1.7	Transition time.....	105
128.7.1.8	Transmit jitter test requirements	105
128.7.1.9	Transmit jitter	105
128.7.2	Receiver characteristics	106
128.7.2.1	Receiver interference tolerance	106
128.7.2.2	Signaling speed range	106
128.7.2.3	AC-coupling.....	106
128.7.2.4	Input signal amplitude	107
128.7.2.5	Differential input return loss.....	107
128.8	Interconnect characteristics.....	107
128.9	Environmental specifications.....	107
128.9.1	General safety	107
128.9.2	Network safety	107
128.9.3	Installation and maintenance guidelines	107
128.9.4	Electromagnetic compatibility	107
128.9.5	Temperature and humidity	107
128.10	Protocol implementation conformance statement (PICS) proforma for Clause 128, Physical Medium Dependent sublayer and baseband medium, type 2.5GBASE-KX	108
128.10.1	Introduction.....	108
128.10.2	Identification	108
128.10.2.1	Implementation identification.....	108
128.10.2.2	Protocol summary	108
128.10.3	Major capabilities/options.....	109
128.10.4	PICS proforma tables for Clause 128, Physical Medium Dependent (PMD) sublayer and baseband medium, type 2.5GBASE-KX.....	109
128.10.4.1	PMD functional specifications.....	109
128.10.4.2	Management functions.....	110
128.10.4.3	Transmitter electrical characteristics	111
128.10.4.4	Receiver electrical characteristics.....	112
128.10.4.5	Environmental and safety specifications	112
129	Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R.....	113
129.1	Overview.....	113
129.1.1	Scope.....	113
129.1.2	Relationship of 5GBASE-R to other standards	113
129.1.3	Summary of 5GBASE-R sublayers	114
129.1.3.1	Physical Coding Sublayer (PCS).....	114
129.1.3.2	Physical Medium Attachment (PMA) sublayer.....	114
129.1.4	Inter-sublayer interfaces	115
129.2	Physical Coding Sublayer (PCS).....	115
129.2.1	Functions within the PCS	115

129.2.2	Notation conventions	116
129.2.3	Transmission order	116
129.2.4	Low Power Idle	117
129.2.5	PCS used with 5GBASE-KR PMD	118
129.3	Physical Medium Attachment (PMA) sublayer	118
129.3.1	Service Interface	118
129.3.2	Functions within the PMA	118
129.3.2.1	PMA transmit function	119
129.3.2.2	PMA receive function	119
129.3.3	PMA loopback mode (optional)	119
129.4	Compatibility considerations	119
129.5	Delay constraints	119
129.6	Environmental specifications	119
129.7	Protocol implementation conformance statement (PICS) proforma for Clause 129, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R	120
129.7.1	Introduction	120
129.7.2	Identification	120
129.7.2.1	Implementation identification	120
129.7.2.2	Protocol summary	120
129.7.3	Major capabilities/options	121
129.7.4	PICS Proforma Tables for PCS, type 5GBASE-R	121
129.7.4.1	Coding rules	121
129.7.4.2	Scrambler and Descrambler	121
129.7.5	Test-pattern modes	122
129.7.5.1	Bit order	122
129.7.6	Management	122
129.7.6.1	State diagrams	123
129.7.6.2	Loopback	123
129.7.6.3	Delay Constraints	123
129.7.6.4	Auto-Negotiation for Backplane Ethernet functions	123
129.7.6.5	LPI functions	124
130	Physical Medium Dependent sublayer and baseband medium, type 5GBASE-KR	125
130.1	Overview	125
130.2	Physical Medium Dependent (PMD) service interface	125
130.2.1	PMD_UNITDATA.request	126
130.2.1.1	Semantics of the service primitive	126
130.2.1.2	When generated	126
130.2.1.3	Effect of receipt	126
130.2.2	PMD_UNITDATA.indication	126
130.2.2.1	Semantics of the service primitive	126
130.2.2.2	When generated	126
130.2.2.3	Effect of receipt	126
130.2.3	PMD_SIGNAL.indication	126
130.2.3.1	Semantics of the service primitive	126
130.2.3.2	When generated	127
130.2.3.3	Effect of receipt	127
130.2.4	PMD_RX_MODE.request	127
130.2.4.1	Semantics of the service primitive	127
130.2.4.2	When generated	127
130.2.4.3	Effect of receipt	127
130.2.5	PMD_TX_MODE.request	127

130.2.5.1	Semantics of the service primitive.....	127
130.2.5.2	When generated	127
130.2.5.3	Effect of receipt	128
130.3	PCS requirements for Auto-Negotiation (AN) service interface.....	128
130.4	Delay constraints.....	128
130.5	PMD MDIO function mapping.....	128
130.6	PMD functional specifications.....	129
130.6.1	Link block diagram.....	129
130.6.2	PMD transmit function	129
130.6.3	PMD receive function.....	129
130.6.4	PMD signal detect function	129
130.6.5	PMD transmit disable function.....	130
130.6.6	Loopback mode.....	130
130.6.7	PMD_fault function	131
130.6.8	PMD transmit fault function.....	131
130.6.9	PMD receive fault function.....	131
130.6.10	PMD LPI function	131
130.7	5GBASE-KR electrical characteristics	132
130.7.1	Transmitter characteristics	132
130.7.1.1	Test fixture.....	132
130.7.1.2	Test fixture characteristics	133
130.7.1.3	Signaling speed	133
130.7.1.4	Output amplitude	133
130.7.1.5	Differential output return loss.....	134
130.7.1.6	Common-mode output return loss	135
130.7.1.7	Transition time.....	135
130.7.1.8	Transmit jitter test requirements.....	135
130.7.1.9	Transmit jitter	136
130.7.1.10	Transmitter output waveform.....	136
130.7.2	Receiver characteristics	137
130.7.2.1	Receiver interference tolerance	138
130.7.2.2	Signaling speed range.....	138
130.7.2.3	AC-coupling.....	138
130.7.2.4	Input signal amplitude	138
130.7.2.5	Differential input return loss.....	138
130.8	Interconnect characteristics.....	139
130.9	Environmental specifications.....	139
130.9.1	General safety	139
130.9.2	Network safety	139
130.9.3	Installation and maintenance guidelines	139
130.9.4	Electromagnetic compatibility	139
130.9.5	Temperature and humidity.....	139
130.10	Protocol implementation conformance statement (PICS) proforma for Clause 130, Physical Medium Dependent (PMD) sublayer and baseband medium, type 5GBASE-KR.....	140
130.10.1	Introduction.....	140
130.10.2	Identification.....	140
130.10.2.1	Implementation identification.....	140
130.10.2.2	Protocol summary	140
130.10.3	Major capabilities/options.....	141
130.10.4	PICS proforma tables for Clause 130, Physical Medium Dependent (PMD) sublayer and baseband medium, type 5GBASE-KR.....	141
130.10.4.1	PCS requirements for AN service interface.....	141
130.10.4.2	PMD functional specifications.....	142
130.10.4.3	Management functions.....	143

130.10.4.4 PMD Transmitter electrical characteristics.....	143
130.10.4.5 Receiver electrical characteristics.....	144
130.10.4.6 Environmental specifications.....	144
Annex 31B (normative) MAC Control PAUSE operation.....	145
31B.3 Detailed specification of PAUSE operation.....	145
31B.3.7 Timing considerations for PAUSE operation.....	145
31B.4 Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation.....	146
31B.4.3 Major capabilities/options.....	146
31B.4.6 PAUSE command MAC timing considerations.....	146
Annex 69A (normative) Interference tolerance testing.....	147
69A.1 Introduction.....	147
69A.2 Test setup.....	147
69A.2.1 Pattern generator.....	147
69A.2.2 Test channel.....	148
69A.2.4 Transmitter control.....	148
69A.3 Test methodology.....	148
Annex 69B (normative) Interconnect characteristics.....	149
69B.4 Channel parameters.....	149
69B.4.1 Overview.....	149
69B.4.2 Fitted attenuation.....	149
69B.4.3 Insertion loss.....	151
69B.4.4 Insertion loss deviation.....	153
69B.4.5 Return loss.....	155
69B.4.6 Crosstalk.....	156
69B.4.6.4 Insertion loss to crosstalk ratio (ICR).....	156
Annex 127A (informative) Compatibility of 2.5GBASE-X PCS/PMA with 1000BASE-X PCS/PMA running 2.5 times faster.....	158
Annex 128A (normative) 2.5 Gb/s Storage Enclosure Interface (2.5GSEI).....	159
128A.1 Overview.....	159
128A.1.1 Bit error ratio.....	161
128A.2 2.5GSEI compliance point definitions.....	161
128A.3 2.5GSEI electrical characteristics.....	164
128A.3.1 2.5GSEI host output characteristics.....	164
128A.3.1.1 Signaling rate and range.....	164
128A.3.1.2 Signaling levels.....	165
128A.3.1.3 Output return loss.....	165
128A.3.1.4 Transmit jitter test requirements.....	166
128A.3.1.5 Transmit jitter.....	166
128A.3.1.6 Transmitter output noise and distortion.....	166
128A.3.2 2.5GSEI host input characteristics.....	167
128A.3.2.1 Input differential return loss.....	167
128A.3.2.2 Receiver interference tolerance.....	167
128A.3.2.3 Receiver jitter tolerance.....	169
128A.3.3 2.5GSEI drive output characteristics.....	170

128A.3.4 2.5GSEI drive input characteristics	171
128A.3.4.1 Input differential return loss	171
128A.3.4.2 Receiver interference tolerance	171
128A.3.4.3 Receiver jitter tolerance	172
128A.4 Protocol implementation conformance statement (PICS) proforma for Annex 128A, 2.5 Gb/s Storage Enclosure Interface (2.5GSEI)	174
128A.4.1 Introduction	174
128A.4.2 Identification	174
128A.4.2.1 Implementation identification	174
128A.4.2.2 Protocol summary	174
128A.4.3 Major capabilities/options	175
128A.4.4 PICS proforma tables for 2.5 Gb/s Storage Enclosure Interface (2.5GSEI)	175
128A.4.4.1 Host output functions	175
128A.4.4.2 Host input functions	175
128A.4.4.3 Drive output functions	176
128A.4.4.4 Drive input functions	176
Annex 128B (normative) Test fixtures for 2.5 Gb/s and 5 Gb/s Storage Enclosure Interfaces	177
128B.1 Host and drive compliance boards	177
128B.1.1 Test fixture return loss	177
128B.1.2 Test fixture insertion loss	177
128B.2 Mated test fixtures	178
128B.2.1 Mated test fixtures insertion loss	179
128B.2.2 Mated test fixtures return loss	179
128B.2.3 Mated test fixtures integrated crosstalk noise	180
128B.2.3.1 Mated test fixture near-end crosstalk (NEXT) loss	180
128B.3 Protocol implementation conformance statement (PICS) proforma for Annex 128B, Test fixtures for 2.5 Gb/s and 5 Gb/s Storage Enclosure Interfaces	182
128B.3.1 Introduction	182
128B.3.2 Identification	182
128B.3.2.1 Implementation identification	182
128B.3.2.2 Protocol summary	182
128B.3.3 Major capabilities/options	183
128B.3.4 PICS proforma tables for test fixtures	183
128B.3.4.1 Management functions	183
Annex 130A (normative) 5 Gb/s Storage Enclosure Interface (5GSEI)	184
130A.1 Overview	184
130A.1.1 Bit error ratio	186
130A.2 5GSEI compliance point definitions	186
130A.3 5GSEI electrical characteristics	189
130A.3.1 5GSEI host output characteristics	189
130A.3.1.1 Signaling rate and range	189
130A.3.1.2 Signaling levels	190
130A.3.1.3 Output return loss	190
130A.3.1.4 Transmitter output waveform	191
130A.3.1.4.1 Linear fit to the measured waveform	191
130A.3.1.4.2 Steady-state voltage and linear fit pulse peak	191
130A.3.1.4.3 Pre-cursor coefficient	191
130A.3.1.5 Transmit jitter test requirements	192
130A.3.1.6 Transmit jitter	192
130A.3.1.7 Transmitter output noise and distortion	192

130A.3.2 5GSEI host input characteristics	193
130A.3.2.1 Input differential return loss	193
130A.3.2.2 Receiver interference tolerance	193
130A.3.2.3 Receiver jitter tolerance	195
130A.3.3 5GSEI drive output characteristics	196
130A.3.3.1 Linear fit to the measured waveform	197
130A.3.3.2 Steady-state voltage and linear fit pulse peak	197
130A.3.3.3 Precursor coefficient	197
130A.3.3.4 Transmitter output noise and distortion	197
130A.3.4 5GSEI drive input characteristics	198
130A.3.4.1 Input differential return loss	198
130A.3.4.2 Receiver interference tolerance	198
130A.3.4.3 Receiver jitter tolerance	200
130A.4 Protocol implementation conformance statement (PICS) proforma for Annex 130A.5 Gb/s Storage Enclosure Interface (5GSEI)	202
130A.4.1 Introduction	202
130A.4.2 Identification	202
130A.4.2.1 Implementation identification	202
130A.4.2.2 Protocol summary	202
130A.4.3 Major capabilities/options	203
130A.4.4 PICS proforma tables for 5 Gb/s Storage Enclosure Interface (5GSEI)	203
130A.4.4.1 Host output functions	203
130A.4.4.2 Host input functions	203
130A.4.4.3 Drive output functions	204
130A.4.4.4 Drive input functions	204

IEEE Standard for Ethernet

Amendment 1: Physical Layer Specifications and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation over Backplane

[This amendment is based on IEEE Std 802.3™-2018.]

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

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

Cross references that refer to clauses, tables, equations, or figures not covered by this amendment are highlighted in green.¹

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

1. Introduction

1.3 Normative references

Insert the following new reference (for SFF-8482 connector) in alphanumerical order:

SFF-8482, Specification for Serial Attachment 2X Unshielded Connector.²

1.4 Definitions

Insert the following new definition (for 2.5GBASE-KX) after 1.4.81 1G-EPON:

1.4.81a 2.5GBASE-KX: IEEE 802.3 Physical Layer specification for 2.5 Gb/s using 2.5GBASE-X encoding over an electrical backplane. (See IEEE Std 802.3, Clause 128.)

Insert the following new definitions (for 2.5GBASE-X, 2.5GP11, and 2.5GSEI) after 1.4.82 2.5GBASE-T:

1.4.82a 2.5GBASE-X: IEEE 802.3 physical coding and physical medium attachment for serial 2.5 Gb/s operation. (See IEEE Std 802.3, Clause 127.)

1.4.82b 2.5GP11: The 2.5 Gb/s PCS Internal Interface is a logical interface that is internal to the 2.5GBASE-X PCS and exists solely for the purposes of defining the 2.5GBASE-X PCS functionality. (See IEEE Std 802.3, 127.2.5.1.)

1.4.82c 2.5GSEI: The 2.5 Gb/s Interface from storage sub-system or bridge subsystem onto backplane (See IEEE Std 802.3, Annex 128A.)

Insert the following new definitions (for 5GBASE-KR and 5GBASE-R) after 1.4.128 4D-PAM5:

1.4.128a 5GBASE-KR: IEEE 802.3 Physical Layer specification for 5 Gb/s using 5GBASE-R encoding over an electrical backplane. (See IEEE Std 802.3, Clause 130.)

1.4.128b 5GBASE-R: An IEEE 802.3 physical coding sublayer and physical medium attachment sublayer for serial 5 Gb/s operation. (See IEEE Std 802.3, Clause 129.)

Insert the following new definition (for 5GSEI) after 1.4.129 5GBASE-T:

1.4.129a 5GSEI: The 5 Gb/s Interface from storage sub-system or bridge subsystem onto backplane (See IEEE Std 802.3, Annex 130A.)

Change 1.4.150 as shown:

1.4.150 BASE-R: An IEEE 802.3 family of Physical Layer devices using 64B/66B encoding defined in Clause 49, Clause 82, Clause 107, or Clause 119, or Clause 129 of IEEE Std 802.3. (See ~~IEEE Std 802.3, Clause 49, Clause 82, Clause 107, Clause 119.~~)

²SFF documents are available from the Storage Networking Industry Association (www.snia.org/sff/specifications).

1.5 Abbreviations

Insert the following new abbreviations into 1.5 in alphanumeric order:

2.5GPII	2.5 Gb/s PCS Internal Interface
2.5GSEI	2.5 Gb/s Storage Enclosure Interface
5GSEI	5 Gb/s Storage Enclosure Interface

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

30. Management

30.3 Layer management for DTEs

30.3.2 PHY device managed object class

30.3.2.1 PHY device attributes

30.3.2.1.2 aPhyType

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.3.2.1.2 after the entry for 2.5GBASE-T:

2.5GBASE-X Clause 127 2.5 Gb/s 8B/10B

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.3.2.1.2 after the entry for 5GBASE-T:

5GBASE-R Clause 129 5 Gb/s 64/66B

30.3.2.1.3 aPhyTypeList

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.3.2.1.3 after the entry for 2.5GBASE-T:

2.5GBASE-X Clause 127 2.5 Gb/s 8B/10B

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.3.2.1.3 after the entry for 5GBASE-T:

5GBASE-R Clause 129 5 Gb/s 64/66B

30.3.2.1.5 aSymbolErrorDuringCarrier

Change the fourth paragraph of "BEHAVIOUR DEFINED AS" in 30.3.2.1.5 as shown:

BEHAVIOUR DEFINED AS:

For operation at 5 Gb/s, 10 Gb/s, 25 Gb/s, 40 Gb/s, 100 Gb/s, 200 Gb/s, and 400 Gb/s, it is a count of the number of times the receiving media is non-idle (the time between the Start of Packet Delimiter and the End of Packet Delimiter as defined by 46.2.5 and 81.2.5) for a period of time equal to or greater than minFrameSize, and during which there was at least one occurrence of an event that causes the PHY to indicate "Receive Error" on the media independent interface (see Table 46-4 and Table 81-4).

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 entry in "APPROPRIATE SYNTAX" in 30.5.1.1.2 before the entry for 2.5GBASE-T:

2.5GBASE-KX 2.5GBASE-X PMD as specified in Clause 128 over an electrical backplane as specified in Clause 128

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.5.1.1.2 after the entry for 2.5GBASE-T:

2.5GBASE-X 2.5GBASE-X PCS/PMA as specified in Clause 127 over undefined PMD

Insert the following new entries in "APPROPRIATE SYNTAX" in 30.5.1.1.2 before the entry for 5GBASE-T:

5GBASE-KR 5GBASE-KR PMD as specified in Clause 130 over an electrical backplane as specified in Clause 130

5GBASE-R 5GBASE-R PCS/PMA as specified in Clause 129 over undefined PMD

Change the second paragraph of "BEHAVIOUR DEFINED AS" in 30.5.1.1.2 as shown:

The enumerations 1000BASE-X, 1000BASE-XHD, 1000BASE-XFD, 2.5GBASE-X, 5GBASE-R, 10GBASE-X, 10GBASE-R, 10GBASE-W, 25GBASE-R, 40GBASE-R, 100GBASE-R, 200GBASE-R, and 400GBASE-R shall only be returned if the underlying PMD type is unknown.;

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 entry in "APPROPRIATE SYNTAX" in 30.6.1.1.5 after the entry for 2.5GBASE-T:

2.5GKX 2.5GBASE-KX as specified in Clause 128

Insert the following new entry in "APPROPRIATE SYNTAX" in 30.6.1.1.5 after the entry for 5GBASE-T:

5GKR 5GBASE-KR as specified in Clause 130

45. Management Data Input/Output (MDIO) Interface

45.2 MDIO Interface Registers

45.2.1 PMA/PMD registers

Change the name of registers 1.160 and 1.161 in Table 45–3 as shown (unchanged rows not shown):

Table 45–3—PMA/PMD registers

Register address	Register name	Subclause
1.160	1000BASE-KX/ <u>2.5GBASE-KX</u> control	45.2.1.97
1.161	1000BASE-KX/ <u>2.5GBASE-KX</u> status	45.2.1.98

45.2.1.1 PMA/PMD control 1 register (Register 1.0)

45.2.1.1.5 PMA local loopback (1.0.0)

Change the first sentence of the second paragraph of 45.2.1.1.5 as shown:

The local loopback function is mandatory for the 1000BASE-KX, 2.5GBASE-X, 5GBASE-R, 10GBASE-KR, 10GBASE-X, 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 port type and optional for all other port types, except 2BASE-TL, 10PASS-TS, and 10/1GBASE-PRX, which do not support loopback.

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

Change the description for bits 1.7.6:0 in Table 45–7 as shown (unchanged table rows and bit description lines not shown):

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

Bit(s)	Name	Description	R/W ^a
1.7.6:0	PMA/PMD type selection	6 5 4 3 2 1 0 0 1 1 1 1 0 0 = <u>5GBASE-KR</u> PMA/PMD reserved 0 1 1 1 0 1 1 = <u>2.5GBASE-KX</u> PMA/PMD reserved	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)

Change Table 45-9 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

Table 45-9—Transmit fault description location

PMA/PMD	Description location
<u>2.5GBASE-KX</u>	<u>128.6.8</u>
2.5GBASE-T, 5GBASE-T	126.4.2.2
<u>5GBASE-KR</u>	<u>130.6.8</u>

45.2.1.7.5 Receive fault (1.8.10)

Change Table 45-10 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

Table 45-10—Receive fault description location

PMA/PMD	Description location
<u>2.5GBASE-KX</u>	<u>128.6.9</u>
2.5GBASE-T, 5GBASE-T	126.4.2.4
<u>5GBASE-KR</u>	<u>130.6.9</u>

45.2.1.8 PMD transmit disable register (Register 1.9)

Change Table 45-12 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

Table 45-12—Transmit disable description location

PMA/PMD	Description location
<u>2.5GBASE-KX</u>	<u>128.6.5</u>
2.5GBASE-T and 5GBASE-T	126.4.2.3
<u>5GBASE-KR</u>	<u>130.6.5</u>

45.2.1.18 2.5G/5G PMA/PMD extended ability register (Register 1.21)

Change Table 45–21 as shown:

Table 45–21— 2.5G/5G PMA/PMD extended ability register bit definitions

Bit(s)	Name	Description	R/W ^a
1.21.15:24	Reserved	Value always 0	RO
1.21.3	5GBASE-KR ability	1 = PMA/PMD is able to perform 5GBASE-KR 0 = PMA/PMD is not able to perform 5GBASE-KR	RO
1.21.2	2.5GBASE-KX ability	1 = PMA/PMD is able to perform 2.5GBASE-KX 0 = PMA/PMD is not able to perform 2.5GBASE-KX	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

Insert the following new subclauses (45.2.1.18.a and 45.2.1.18.b) after the introductory text of 45.2.1.18 (i.e., before 45.2.1.18.1):

45.2.1.18.a 5GBASE-KR ability (1.21.3)

When read as a one, bit 1.21.3 indicates that the PMA/PMD is able to operate as a 5GBASE-KR PMA/PMD type. When read as a zero, bit 1.21.3 indicates that the PMA/PMD is not able to operate as a 5GBASE-KR PMA/PMD type.

45.2.1.18.b 2.5GBASE-KX ability (1.21.2)

When read as a one, bit 1.21.2 indicates that the PMA/PMD is able to operate as a 2.5GBASE-KX PMA/PMD type. When read as a zero, bit 1.21.2 indicates that the PMA/PMD is not able to operate as a 2.5GBASE-KX PMA/PMD type.

45.2.1.89 BASE-R PMD control register (Register 1.150)

Insert the following new paragraph after the first paragraph of 45.2.1.89:

The BASE-R PMD control register is also used by 5GBASE-KR described in Clause 130 to disable the transmitter equalizer for test purposes. 5GBASE-KR does not use the start-up protocol.

Change Table 45-69 as shown:

Table 45–69— BASE-R PMD control register

Bit(s)	Name	Description	R/W ^a
1.150.15:2 1.150.15:3	Reserved	Value always 0	RO
1.150.2	<u>Transmitter equalizer disable</u>	<u>1 = Disable the 5GBASE-KR transmitter equalizer</u> <u>0 = Normal operation</u>	R/W
1.150.1	Training enable	1 = Enable the BASE-R start-up protocol 0 = Disable the BASE-R start-up protocol	R/W
1.150.0	Restart training	1 = Restart BASE-R start-up protocol 0 = Normal operation	R/W SC

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

Insert the following new subclause (45.2.1.89.3) after 45.2.1.89.2:

45.2.1.89.3 Transmitter equalizer disable (1.150.2)

When bit 1.150.2 is set to one, 5GBASE-KR transmitter equalization is disabled. The default value of bit 1.150.2 is zero.

Change the title and text of 45.2.1.97 as shown:

45.2.1.97 1000BASE-KX/2.5GBASE-KX control register (Register 1.160)

The assignment of bits in the 1000BASE-KX/2.5GBASE-KX control register is shown in Table 45–77.

Change the title of Table 45–77 as shown:

Table 45–77—1000BASE-KX/2.5GBASE-KX control register

45.2.1.97.1 PMD transmit disable (1.160.0)

Change 45.2.1.97.1 as shown:

This bit disables the 1000BASE-KX transmitter as defined in 70.6.5 or the 2.5GBASE-KX transmitter as defined in 128.6.5.

Change the title and text of 45.2.1.98 as shown:

45.2.1.98 1000BASE-KX/2.5GBASE-KX status register (Register 1.161)

The assignment of bits in the 1000BASE-KX/2.5GBASE-KX status register is shown in Table 45–78.

Change the title of Table 45–78 as shown:

Table 45–78—1000BASE-KX/2.5GBASE-KX status register

Change the title and text of 45.2.1.98.6 as shown:

45.2.1.98.6 1000BASE-KX/2.5GBASE-KX signal detect (1.161.0)

The PMD signal detect function is optional for both 1000BASE-KX PMD (see 70.6.4) and 2.5GBASE-KX PMD (see 128.6.4) is mandatory if EEE is implemented and optional otherwise. The 1000BASE-X PCS and 2.5GBASE-X PCS requires signal detect to be one before synchronization can occur. If the signal detect function is not implemented, this bit is set to one.

45.2.1.140 Test-pattern ability (Register 1.1500)

Change the first paragraph of 45.2.1.140 as shown:

The test-pattern ability register is used for PHY types that implement SSPRQ, square wave, and PRBS testing in the PMA. These functions are described in 127.3.4.1, 83.5.10, and 120.5.11. The assignment of bits in the test-pattern ability register is shown in Table 45–105.

45.2.3 PCS registers

Change the indicated rows of Table 45–176 as shown (unchanged rows not shown):

Table 45–176—PCS registers

Register address	Register name	Subclause
3.34 through 3.37	5/10/25GBASE-R PCS test pattern seed A	45.2.3.17
3.38 through 3.41	5/10/25GBASE-R PCS test pattern seed B	45.2.3.18

45.2.3.1 PCS control 1 register (Register 3.0)

45.2.3.1.2 Loopback (3.0.14)

Change the first paragraph of 45.2.3.1.2 as shown:

When the 100BASE-T1, any MultiGBASE-T, or the 5/10GBASE-R mode of operation is selected for the PCS using the PCS type selection field (3.7.3:0), 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 100BASE-T1, 5/10GBASE-R, 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 100BASE-T1 PCS during loopback is specified in 96.3.5. The specific behavior of the 5/10GBASE-R PCS during loopback is specified in 49.2. The specific behavior for the 10GBASE-T PCS during loopback is specified in 55.3.7.3. 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.

45.2.3.6 PCS control 2 register (Register 3.7)

Change the description for bits 3.7.3:0 in Table 45–180 as shown (unchanged table rows and bit description lines not shown):

Table 45–180—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 1 1 x= reserved 1 1 1 1= Select 5GBASE-R PCS type 1 1 1 0= Select 2.5GBASE-X PCS type	R/W

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

45.2.3.8 PCS status 3 register (Register 3.9)

Change Table 45–182 as shown:

Table 45–182—PCS status 3 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.9.15:24	Reserved	Value always 0	RO
3.9.3	5GBASE-R capable	1 = PCS is able to support 5GBASE-R PCS type 0 = PCS is not able to support 5GBASE-R PCS type	RO
3.9.2	2.5GBASE-X capable	1 = PCS is able to support 2.5GBASE-X PCS type 0 = PCS is not able to support 2.5GBASE-X PCS type	RO
3.9.1	400GBASE-R capable	1 = PCS is able to support 400GBASE-R PCS type 0 = PCS is not able to support 400GBASE-R PCS type	RO
3.9.0	200GBASE-R capable	1 = PCS is able to support 200GBASE-R PCS types 0 = PCS is not able to support 200GBASE-R PCS types	RO

^aRO = Read only

Insert the following new subclauses (45.2.3.8.a and 45.2.3.8.b) after the introductory text of 45.2.3.8 (i.e., before 45.2.3.8.1):

45.2.3.8.a 5GBASE-R capable (3.9.3)

When read as a one, bit 3.9.3 indicates that the PCS is able to support the 5GBASE-R PCS type. When read as a zero, bit 3.9.3 indicates that the PCS is not able to support the 5GBASE-R PCS type.

45.2.3.8.b 2.5GBASE-X capable (3.9.2)

When read as a one, bit 3.9.2 indicates that the PCS is able to support operation in a 2.5GBASE-X PCS type. When read as a zero, bit 3.9.2 indicates that the PCS is not able to support operation in a 2.5GBASE-X PCS type.

45.2.3.11 EEE control and capability 2 register (Register 3.21)

Change Table 45–184 as shown:

Table 45–184—EEE control and capability 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
3.21.15:69	Reserved	Value always 0	RO
<u>3.21.8</u>	<u>5GBASE-KR EEE</u>	<u>1 = EEE is supported for 5GBASE-R</u> <u>0 = EEE is not supported for 5GBASE-R</u>	<u>RO</u>
<u>3.21.7</u>	<u>2.5GBASE-KX EEE</u>	<u>1 = EEE is supported for 2.5GBASE-X</u> <u>0 = EEE is not supported for 2.5GBASE-X</u>	<u>RO</u>
<u>3.21.6</u>	<u>Reserved</u>	<u>Value always 0</u>	<u>RO</u>
3.21.5	400GBASE-R fast wake	1 = EEE fast wake is supported for 400GBASE-R 0 = EEE fast wake is not supported for 400GBASE-R	RO
3.21.4	Reserved	Value always 0	RO
3.21.3	200GBASE-R fast wake	1 = EEE fast wake is supported for 200GBASE-R 0 = EEE fast wake is not supported for 200GBASE-R	RO
3.21.2	25GBASE-T EEE	1 = EEE is supported for 25GBASE-T 0 = EEE is not supported for 25GBASE-T	RO
3.21.1	5GBASE-T EEE	1 = EEE is supported for 5GBASE-T 0 = EEE is not supported for 5GBASE-T	RO
3.21.0	2.5GBASE-T EEE	1 = EEE is supported for 2.5GBASE-T 0 = EEE is not supported for 2.5GBASE-T	RO

^aRO = Read only

Insert the following new subclauses (45.2.3.11.a and 45.2.3.11.b) after the introductory text of 45.2.3.11 (i.e., before 45.2.3.11.1):

45.2.3.11.a 5GBASE-KR EEE supported (3.21.8)

If the device supports EEE operation for 5GBASE-KR as defined in 130.1, this bit shall be set to one.

45.2.3.11.b 2.5GBASE-KX EEE supported (3.21.7)

If the device supports EEE operation for 2.5GBASE-KX as defined in 128.1, this bit shall be set to one.

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

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

Change last sentence of 45.2.3.15.1 as shown:

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

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

Change the last sentence of the first paragraph of 45.2.3.15.4 as shown:

This bit is a direct reflection of the state of the hi_ber variable in the BER monitor state diagrams as defined in 49.2.13.2.2 for 5/10/25GBASE-R and in 82.2.19.2.2 for 40/100GBASE-R.

45.2.3.15.5 BASE-R and MultiGBASE-T PCS block lock (3.32.0)

Change the third sentence of 45.2.3.15.5 as shown:

For a 10GBASE-R or 5/10/25GBASE-R PCS, this bit is a direct reflection of the state of the block_lock variable defined in 49.2.13.2.2.

45.2.3.16 BASE-R and MultiGBASE-T PCS status 2 register (Register 3.33)

45.2.3.16.1 Latched block lock (3.33.15)

Change 45.2.3.16.1 as shown:

When read as a one, bit 3.33.15 indicates that the 5/10/25/40/100GBASE-R or a member of the MultiGBASE-T set PCS has achieved block lock. When read as a zero, bit 3.33.15 indicates that the 5/10/25/40/100GBASE-R or a member of the MultiGBASE-T set PCS has lost block lock.

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

This bit is a latching low version of the 5/10/25/40/100GBASE-R and MultiGBASE-T PCS block lock status bit (3.32.0).

45.2.3.16.2 Latched high BER (3.33.14)

Change 45.2.3.16.2 as shown:

When read as a one, bit 3.33.14 indicates that the 5/10/25/40/100GBASE-R or a member of the MultiGBASE-T set PCS has detected a high BER. When read as a zero, bit 3.33.14 indicates that the 5/10/25/40/100GBASE-R or a member of the MultiGBASE-T set PCS has not detected a high BER.

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

This bit is a latching high version of the 5/10/25/40/100GBASE-R and MultiGBASE-T PCS high BER status bit (3.32.1).

45.2.3.16.3 BER(3.33.13:8)

Change the first sentence of 45.2.3.16.3 as shown:

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 5/10/25/40/100GBASE-R and defined by counter lfer_count in 126.3.7.2 for 2.5GBASE-T and 5GBASE-T, in 55.3.7.2 for 10GBASE-T, and in 113.3.7.2 for 25GBASE-T and 40GBASE-T.

45.2.3.16.4 Errored blocks (3.33.7:0)

Change the first sentence of 45.2.3.16.4 as shown:

The errored blocks counter is an eight-bit count defined by the counter errored_block_count specified in 49.2.14.2 for 5/10/25GBASE-R, in 82.3.1 for 40/100GBASE-R, in 126.3.7.2 for 2.5GBASE-T and 5GBASE-T, in 55.3.7.2 for 10GBASE-T, and in 113.3.7.2 for 25GBASE-T and 40GBASE-T.

Change the title and the first and second sentences of 45.2.3.17 as shown:

45.2.3.17 5/10/25GBASE-R PCS test pattern seed A (Registers 3.34 through 3.37)

The assignment of bits in the 5/10/25GBASE-R PCS test pattern seed A registers is shown in Table 45–189. This register is only required when the 5GBASE-R, 10GBASE-R, or 25GBASE-R capability is supported.

Change the title of Table 45–189 as shown:

Table 45–189—5/10/25GBASE-R PCS test pattern seed A 0-3 register bit definitions

Change the title and the first and second sentences of 45.2.3.18 as shown:

45.2.3.18 5/10/25GBASE-R PCS test pattern seed B (Registers 3.38 through 3.41)

The assignment of bits in the 5/10/25GBASE-R PCS test pattern seed B registers is shown in Table 45–190. This register is only required when the 5GBASE-R, 10GBASE-R, or 25GBASE-R capability is supported.

Change the title of Table 45–190 as shown:

Table 45–190—5/10/25GBASE-R PCS test pattern seed B 0-3 register bit definitions

45.2.7 Auto-Negotiation registers

45.2.7.2 AN status (Register 7.1)

45.2.7.2.1 Parallel detection fault (7.1.9)

Change 45.2.7.2.1 as shown:

The parallel detection Fault bit (7.1.9) shall be set to one to indicate that more than one of 1000BASE-KX, 2.5GBASE-KX, or 10GBASE-KX4 PMAs have indicated link_status=OK when the autoneg_wait_timer expires. The parallel detection fault bit shall be reset to zero on a read of the AN status register (Register 7.1).

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

Change Table 45–318 by adding the following new rows for 7.48.15 and 7.48.14 in numeric order and deleting the Reserved row as shown (unchanged rows not shown):

Table 45–318—Backplane Ethernet, BASE-R copper status register (Register 7.48) bit definitions

Bit(s)	Name	Description	R/W ^a
7.48.15:14	Reserved	Value always 0	RO
7.48.15	5GBASE-KR	1 = PMA/PMD is negotiated to perform 5GBASE-KR 0 = PMA/PMD is not negotiated to perform 5GBASE-KR	RO
7.48.14	2.5GBASE-KX	1 = PMA/PMD is negotiated to perform 2.5GGBASE-KX 0 = PMA/PMD is not negotiated to perform 2.5GGBASE-KX	RO

^aRO = Read only

Change the title of 45.2.7.12.3 as shown:

45.2.7.12.3 Negotiated Port Type (7.48.1, 7.48.2, 7.48.3, 7.48.5, 7.48.6, 7.48.8, 7.48.9, 7.48.10, 7.48.11, 7.48.12, 7.48.13, 7.48.14, 7.48.15)

45.2.7.15 EEE advertisement 2 (Register 7.62)

Change Table 45–321 as shown:

Table 45–321—EEE advertisement 2 register (Register 7.62) bit definitions

Bit(s)	Name	Description	Subclause reference; Next Page bit number	R/W ^a
7.62.15:24	Reserved	Value always 0		RO
7.62.3	5GBASE-KR EEE	1 = Advertise that the 5GBASE-KR has EEE capability 0 = Do not advertise that the 5GBASE-KR has EEE capability	73.7.7.1; U17	R/W
7.62.2	2.5GBASE-KX EEE	1 = Advertise that the 2.5GBASE-KX has EEE capability 0 = Do not advertise that the 2.5GBASE-KX has EEE capability	73.7.7.1; U16	R/W
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

Insert the following new subclauses (45.2.7.15.a and 45.2.7.15.b) after the introductory text of 45.2.7.15 (i.e., before 45.2.7.15.1):

45.2.7.15.a 5GBASE-KR EEE (7.62.3)

If the device supports EEE operation for 5GBASE-KR as defined in 130.1 and EEE operation is desired, this bit shall be set to one.

45.2.7.15.b 2.5GBASE-KX EEE (7.62.2)

If the device supports EEE operation for 2.5GBASE-KX as defined in 128.1 and EEE operation is desired, this bit shall be set to one.

45.2.7.16 EEE link partner ability 2 (Register 7.63)

Change Table 45–322 as shown:

Table 45–322—EEE advertisement 2 register (Register 7.63) bit definitions

Bit(s)	Name	Description	Subclause reference; Next Page bit number	R/W ^a
7.62.15:24	Reserved	Value always 0		RO
7.63.3	5GBASE-KR EEE	1 = Link partner is advertising that the 5GBASE-KR has EEE capability 0 = Link partner is not advertising that the 5GBASE-KR has EEE capability	73.7.7.1; U17	RO
7.63.2	2.5GBASE-KX EEE	1 = Link partner is advertising that the 2.5GBASE-KX has EEE capability 0 = Link partner is not advertising that the 2.5GBASE-KX has EEE capability	73.7.7.1; U16	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

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

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

45.5.3.2 PMA/PMD MMD options

Change the table in 45.5.3.2 by adding the following new rows for 2.5KX and 5KR as shown (other unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
*2.5T	Implementation of a 2.5GBASE-T PMA	45.2.1.18		PMA:O	Yes [] No []
<u>2.5GKX</u>	<u>Implementation of 2.5GBASE-KX PMA/PMD</u>	<u>45.2.1.18.b</u>		<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
<u>5GKR</u>	<u>Implementation of 5GBASE-KR PMA/PMD</u>	<u>45.2.1.18.a</u>		<u>PMA:O</u>	<u>Yes []</u> <u>No []</u>
*5T	Implementation of a 2.5GBASE-T PMA	45.2.1.18		PMA:O	Yes [] No []

45.5.3.3 PMA/PMD management functions

Change the MM124 row in the table in 45.5.3.3 as shown (unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
MM124	Backplane Ethernet AN ability bit is set to one for 1000BASE-KX, 2.5GBASE-KX, 5GBASE-KR, 10GBASE-KX4, and 10GBASE-KR PHYs.	45.2.7.12.4		<u>2.5GKX:M</u> <u>5GKR:M</u> KX:M KX4:M KR:M	Yes [] N/A []

45.5.3.6 PCS options

Insert the following new rows (for 2.5GX and 5GX) at the end of the table in 45.5.3.6:

Item	Feature	Subclause	Value/Comment	Status	Support
2.5GX	Implementation of 2.5GBASE-X PCS	45.2.3.8		PCS:O	Yes [] No [] N/A []
5GR	Implementation of 5GBASE-R PCS	45.2.3.8		PCS:O	Yes [] No [] 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.

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

46.3 XGMII functional specifications

46.3.3 Error and fault handling

46.3.3.3 Response to received invalid frame sequences

Insert the following new text at the end of 46.3.3.3:

To support 2.5GBASE-X compatibility with a 1000BASE-X PCS/PMA running 2.5 times faster as described in Annex 127A, a 2.5 Gb/s MAC/RS implementation connected to a 2.5GBASE-X PHY is required to support an SFD received on either lane 2 or lane 3.

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

69. Introduction to Ethernet operation over electrical backplanes

69.1 Overview

69.1.1 Scope

Change the first sentence of the second paragraph of 69.1.1 as shown:

Backplane Ethernet supports the IEEE 802.3 full duplex MAC operating at 1000 Mb/s, 2.5 Gb/s, 5 Gb/s, 10 Gb/s, 25 Gb/s, 40 Gb/s, or 100 Gb/s providing a bit error ratio (BER) better than or equal to 10^{-12} at the MAC/PLS service interface.

Insert the following new lines after the 1000BASE-KX line in the list after the second paragraph of 69.1.1:

- 2.5GBASE-KX for 2.5 Gb/s operation over a single lane
- 5GBASE-KR for 5 Gb/s operation over a single lane

69.1.2 Relationship of Backplane Ethernet to the ISO OSI reference model

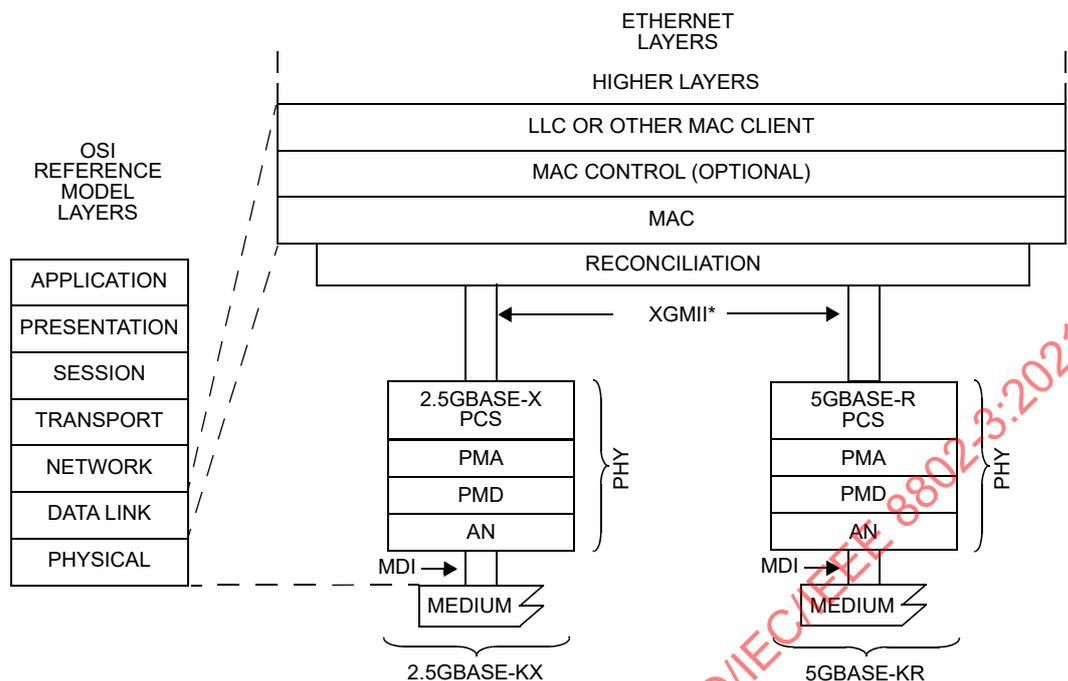
Change the first paragraph of 69.1.2 as shown:

Backplane Ethernet couples the IEEE 802.3 MAC to a family of Physical Layers defined for operation over electrical backplanes. The relationships among Backplane Ethernet, the IEEE 802.3 MAC, and the ISO Open System Interconnection (OSI) reference model are shown in [Figure 69-1](#), [Figure 69-2](#), and [Figure 69-3](#), and [Figure 69-4](#).

Insert the new figure on the next page (Figure 69-4) after Figure 69-3.

Change list item i) in the lettered list of 69.1.2 as shown:

- i) The MDIs for 1000BASE-KX, 2.5GBASE-KX, 5GBASE-KR, 10GBASE-KR, 25GBASE-KR, and 25GBASE-KR-S use a serial data path while the MDIs for 10GBASE-KX4, 40GBASE-KR4, 100GBASE-KR4, and 100GBASE-KP4 use a four-lane data path.



AN = AUTO-NEGOTIATION
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE

*XGMII IS OPTIONAL

Figure 69-4—Architectural positioning of 2.5 Gb/s and 5 Gb/s Backplane Ethernet

69.2 Summary of Backplane Ethernet Sublayers

69.2.3 Physical Layer signaling systems

Insert the following new paragraphs after the first paragraph (“Backplane Ethernet extends”) in 69.2.3:

Backplane Ethernet also extends the family of 2.5GBASE-X Physical Layer signaling systems to include 2.5GBASE-KX. This embodiment specifies operation at 2.5 Gb/s over two differential, controlled impedance pairs of traces (one pair for transmit, one pair for receive). This system employs the 2.5GBASE-X PCS and PMA as defined in Clause 127. The 2.5GBASE-KX PMD is defined in Clause 128.

Backplane Ethernet also extends the family of 5GBASE-R Physical Layer signaling systems to include the 5GBASE-KR. This embodiment specifies 5 Gb/s operation over two differential, controlled impedance pairs of traces (one pair for transmit, one pair for receive). This system employs the 5GBASE-R PCS and PMA as defined in Clause 129. The 5GBASE-KR PMD is defined in Clause 130.

Change the last paragraph in 69.2.3 as shown:

Table 69-1, Table 69-1a, Table 69-2, and Table 69-3 specify the correlation between nomenclature and clauses. A complete implementation conforming to one or more nomenclatures meets the requirements of the corresponding clauses.

Insert the following new table (Table 69–1a) after Table 69-1:

Table 69–1a—Nomenclature and clause correlation for 2.5 Gb/s and 5 Gb/s Backplane Ethernet Physical Layers

Nomenclature	Clause							
	46		73	78	127	128	129	130
	RS	XGMII	Auto-Negotiation	Energy-Efficient Ethernet (EEE)	2.5GBASE-X PCS/PMA	2.5GBASE-KX PMD	5GBASE-R PCS/PMA	5GBASE-KR PMD
2.5GBASE-KX	M ^a	O ^a	M	O	M	M		
5GBASE-KR	M	O	M	O			M	M

^aO = Optional, M = Mandatory

69.3 Delay constraints

Insert the following new paragraph after the third paragraph (“Table 69-5”) of 69.3:

For 2.5GBASE-KX, normative delay specifications may be found in 127.5 and 128.4. For 5GBASE-KR, normative delay specifications may be found in 129.5 and 130.4. They are also referenced in 125.3.

69.5 Protocol implementation conformance statement (PICS) proforma

Change the first paragraph of 69.5 as shown:

The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 70 through Clause 74, Clause 84, Clause 91, Clause 93, Clause 94, Clause 108, Clause 111, Clause 128, Clause 130, and related annexes demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

70. Physical Medium Dependent sublayer and baseband medium, type 1000BASE-KX

70.5 PMD MDIO function mapping

Change Table 70–2 as shown:

Table 70–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/ bit number	PMD control variable
Reset	Control register 1	1.0.15	PMD_reset
PMD Transmit Disable	1000BASE-KX/2.5GBASE-KX control register	1.160.0	PMD_transmit_disable

Change Table 70–3 as shown:

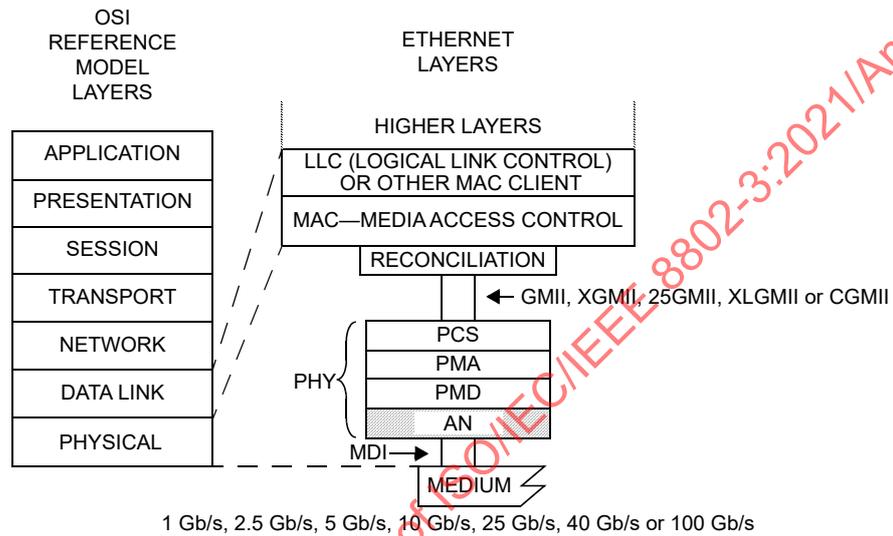
Table 70–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/ bit number	PMD status variable
Fault	Status register 1	1.1.7	PMD_fault
Transmit fault ability	1000BASE-KX/2.5GBASE-KX status register	1.161.13	PMD_Transmit_fault_ability
Receive fault ability	1000BASE-KX/2.5GBASE-KX status register	1.161.12	PMD_Receive_fault_ability
Transmit fault	1000BASE-KX/2.5GBASE-KX status register	1.161.11	PMD_transmit_fault
Receive fault	1000BASE-KX/2.5GBASE-KX status register	1.161.10	PMD_receive_fault
PMD transmit disable ability	1000BASE-KX/2.5GBASE-KX status register	1.161.8	PMD_transmit_disable_ability
Signal detect from PMD	1000BASE-KX/2.5GBASE-KX status register	1.161.0	PMD_signal_detect

73. Auto-Negotiation for backplane and copper cable assembly

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

Replace Figure 73–1 with the following figure (which adds 2.5 Gb/s and 5 Gb/s in the medium rates):



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

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

73.6 Link codeword encoding

73.6.4 Technology Ability Field

Change Table 73–4 by adding the following new rows for A11 and A12 in numeric order and revising the reserved row as shown (unchanged rows not shown):

Table 73–4—Technology Ability Field encoding

Bit	Technology
A11	2.5GBASE-KX
A12	5GBASE-KR
A11 A13 through A22	Reserved for future technology

Change the last paragraph of 73.6.4 as shown:

The fields A[22:13+4] are reserved for future use. Reserved fields shall be sent as zero and ignored on receive.

73.7 Receive function requirements

73.7.4 Arbitration function requirements

73.7.4.1 Parallel Detection function

Change the first two paragraphs of 73.7.4.1 as shown:

The local device detects a link partner that supports Auto-Negotiation by DME page detection. The Parallel Detection function allows detection of link partners that support 1000BASE-KX, 2.5GBASE-KX, and 10GBASE-KX4 but have disabled Auto-Negotiation, and detection of legacy devices that can interoperate with 1000BASE-KX, 2.5GBASE-KX, and 10GBASE-KX4 devices but do not provide [Clause 73](#) Auto-Negotiation.

A local device shall provide Parallel Detection for 1000BASE-KX, 2.5GBASE-KX, and 10GBASE-KX4 if it supports those PHYs. Additionally, parallel detection may be used for 10GBASE-CX4. Parallel detection of 10GBASE-CX4 will be indicated by the setting of the Negotiated Port Type to “10GBASE-KX4 or 10GBASE-CX4” in the management register bit 7.48.2. The means to distinguish between 10GBASE-KX4 and 10GBASE-CX4 is implementation dependent. Parallel Detection shall be performed by directing the MDI receive activity to the PHY. This detection may be done in sequence between detection of DME pages and detection of each supported PHY. If at least one of the 1000BASE-KX, 2.5GBASE-KX, or 10GBASE-KX4 establishes link_status=OK, the LINK STATUS CHECK state is entered and the autoneg_wait_timer is started. If exactly one link_status=OK indication is present when the autoneg_wait_timer expires, then Auto-Negotiation shall set link_control=ENABLE for the PHY indicating link_status=OK. If a PHY is enabled, the Arbitration function shall set link_control=DISABLE to all other PHYs and indicate that Auto-Negotiation has completed. On transition to the AN GOOD CHECK state from the LINK STATUS CHECK state, the Parallel Detection function shall set the bit in the AN LP Base Page ability registers (see [45.2.7.7](#)) corresponding to the technology detected by the Parallel Detection function.

73.7.6 Priority Resolution function

Change Table 73-5 by adding the following new rows for Priority 11 and Priority 12 in numeric order and revising the last row as shown (unchanged rows not shown):

Table 73–5—Priority Resolution

Priority	Technology	Capability
11	<u>5GBASE-KR</u>	<u>5 Gb/s 1 lane</u>
12	<u>2.5GBASE-KX</u>	<u>2.5 Gb/s 1 lane</u>
13	1000BASE-KX	1 Gb/s 1 lane, lowest priority

73.10 State diagrams and variable definitions

73.10.1 State diagram variables

Insert the following new entries (for 2.5GKX and 5GKR) after the entry for 1GKX in the variable list after the first paragraph (‘A variable with “_x]”’) in 73.10.1:

- 2.5GKX; represents the 2.5GBASE-KX PMD.
- 5GKR; represents the 5GBASE-KR PMD.

Change the PD entry in the variable list after the first paragraph (‘A variable with “_x]”’) in 73.10.1 as shown:

- PD; represents all of the following that are present: 1000BASE-KX PMD, 2.5GBASE-KX PMD, and 10GBASE-KX4 (or 10GBASE-CX4) PMD.

Change the single_link_ready entry in the list after the third paragraph (‘Variables of the form “mr_x”,’) in 73.10.1 as shown:

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

- 1) link_status_[1GKX] = OK
- 2) link_status_[2.5GKX] = OK
- 3) link_status_[5GKR] = OK
- 4) link_status_[10GKX4] = OK
- 5) link_status_[10GKR] = OK
- 6) link_status_[25GR] = OK
- 7) link_status_[40GKR4] = OK
- 8) link_status_[40GCR4] = OK
- 9) link_status_[100GCR10] = OK
- 10) link_status_[100GKP4] = OK
- 11) link_status_[100GKR4] = OK
- 12) link_status_[100GCR4] = OK

73.10.2 State diagram timers

Change the two link_fail_inhibit_timer rows in Table 73-7 as shown (unchanged rows not shown):

Table 73-7—Timer min/max value summary

Parameter	Min	Value and tolerance	Max	Units
link_fail_inhibit_timer (when the link is not <u>neither</u> 1000BASE-KX, 2.5GBASE-KX, 5GBASE-KR, or not 10GBASE-KX4)	500		510	ms
link_fail_inhibit_timer (when the link is 1000BASE-KX, <u>2.5GBASE-KX</u> , 5GBASE-KR, or 10GBASE-KX4)	40		50	ms

73.11 Protocol implementation conformance statement (PICS) proforma for Clause 73, Auto-Negotiation for backplane and copper cable assembly⁴

73.11.4 PICS proforma tables for Auto-Negotiation for backplane and copper cable assembly

73.11.4.4 Receive function requirements

Change the RF5 row in the table in 73.11.4.4 as shown:

Item	Feature	Subclause	Value/Comment	Status	Support
RF5	Parallel detection for 1000BASE-KX ₁ , 2.5GBASE-KX , and 10GBASE-KX4	73.7.4.1	Device provides parallel detection if it supports those PHYs	M	Yes []

⁴Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

78.1.1 LPI Signaling

Change the fourth paragraph of 78.1.1 as shown:

The EEE request signals from the PCS control transitions between quiescent and normal operation. The Clause 49 PCS, Clause 107 PCS, and Clause 82 PCS, and Clause 129 PCS also request transmit alert operation to assist the partner device PMD to detect the end of the quiescent state. Additionally ~~these the~~ Clause 49 PCS, Clause 107 PCS, and Clause 82 PCS types generate the RX_LPI_ACTIVE signal which indicates to the Clause 74 BASE-R FEC that it can use rapid block lock because the link partner PCS has bypassed scrambling.

78.1.4 PHY types optionally supporting EEE

Change Table 78-1 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

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

PHY or interface type	Clause
<u>2.5GBASE-KX</u>	<u>127, 128</u>
2.5GBASE-T	126
<u>5GBASE-KR</u>	<u>129, 130</u>

78.2 LPI mode timing parameters description

Change Table 78-2 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

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

PHY or interface type	T_s (μs)		T_q (μs)		T_r (μs)	
	Min	Max	Min	Max	Min	Max
<u>2.5GBASE-KX</u>	<u>19.9</u>	<u>20.1</u>	<u>2 500</u>	<u>2 600</u>	<u>19.9</u>	<u>20.1</u>
2.5GBASE-T	11.52	12.8	76.8	76.8	5.12	5.12
<u>5GBASE-KR</u>	<u>4.9</u>	<u>5.1</u>	<u>1 700</u>	<u>1 800</u>	<u>16.9</u>	<u>17.5</u>

78.5 Communication link access latency

Change Table 78-4 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown (other unchanged rows not shown):

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

PHY or interface type	Case	$T_{w_sys_tx}$ (min) (μ s)	T_{w_phy} (min) (μ s)	$T_{phy_shrink_tx}$ (max) (μ s)	$T_{phy_shrink_rx}$ (max) (μ s)	$T_{w_sys_rx}$ (min) (μ s)
<u>2.5GBASE-KX</u>		<u>13.26</u>	<u>11.25</u>	<u>5.0</u>	<u>6.5</u>	<u>1.76</u>
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
<u>5GBASE-KR</u>		<u>15.38</u>	<u>12.25</u>	<u>5.0</u>	<u>7.5</u>	<u>2.88</u>

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

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

125.1 Overview

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

Change 125.1.2 by adding the following new list item for 2.5GBASE-KX and 5GBASE-KR MDI as shown:

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.
- d) The MDI as specified in Clause 128 for 2.5GBASE-KX and Clause 130 for 5GBASE-KR uses a single-lane data path.

125.1.3 Nomenclature

Change the text of 125.1.3 as shown:

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.

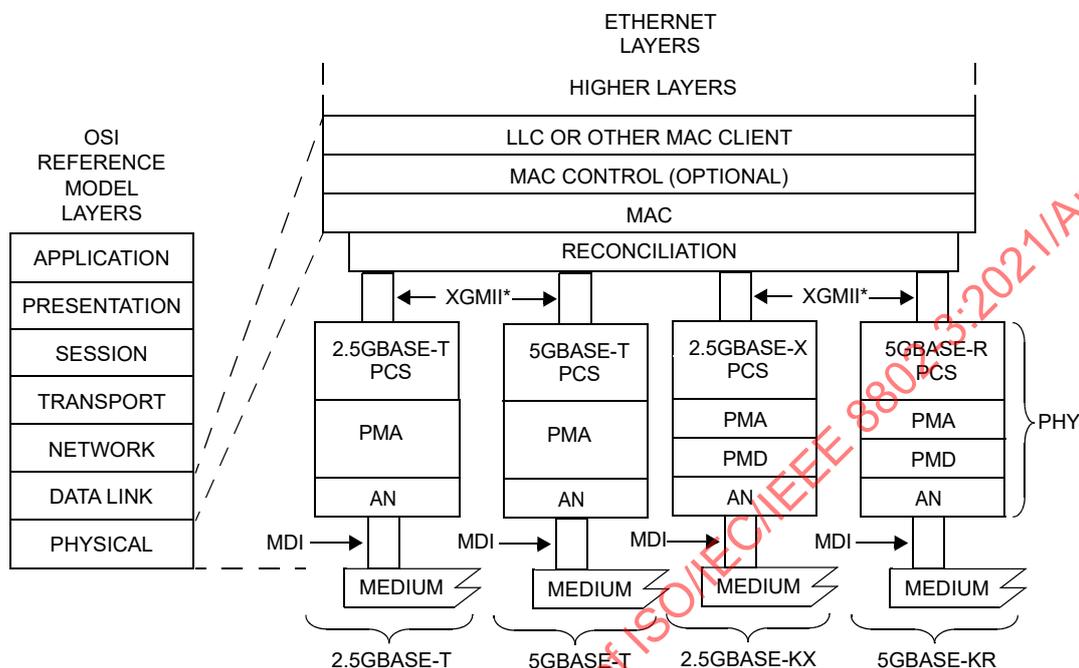
2.5GBASE-T represents Physical Layer devices using Clause 126 Physical Coding Sublayer (PCS), Physical Medium Attachment (PMA) sublayer, and baseband medium, for transmitting 2.5 Gb/s Ethernet over a point-to-point 4-pair balanced twisted-pair medium. 2.5GBASE-T uses low density parity check (LDPC) FEC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

The term 2.5GBASE-X refers to a specific family of Physical Layer implementations based upon the 8B/10B data coding method specified in Clause 127. The 2.5GBASE-X family is composed of 2.5GBASE-KX.

The term 5GBASE-R refers to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 129. The 5GBASE-R family is composed of 5GBASE-KR.

Replace Figure 125-1 with the following figure (which adds the 2.5GBASE-KX and 5GBASE-KR related layer diagram):



AN = AUTO-NEGOTIATION
LLC = LOGICAL LINK CONTROL
MAC = MEDIA ACCESS CONTROL
MDI = MEDIUM DEPENDENT INTERFACE
PCS = PHYSICAL CODING SUBLAYER

PHY = PHYSICAL LAYER DEVICE
PMA = PHYSICAL MEDIUM ATTACHMENT
PMD = PHYSICAL MEDIUM DEPENDENT
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

Change Table 125-1 by adding the following new rows for 2.5GBASE-KX and 5GBASE-KR as shown:

Table 125-1—2.5 Gb/s and 5 Gb/s PHYs

Name	Description
<u>2.5GBASE-KX</u>	<u>2.5 Gb/s PHY using 2.5GBASE-X encoding over one lane of an electrical backplane (see Clause 128)</u>
2.5GBASE-T	2.5 Gb/s PHY using LDPC encoding and PAM16 modulation over balanced twisted-pair structured cabling systems (see Clause 126)
<u>5GBASE-KR</u>	<u>5 Gb/s PHY using 5GBASE-R encoding over one lane of an electrical backplane (see Clause 130)</u>
5GBASE-T	5 Gb/s PHY using LDPC encoding and PAM16 modulation over balanced twisted-pair structured cabling systems (see Clause 126)

Change Table 125-2 by adding the following new columns for clause references and new rows for 2.5GBASE-KX and 5GBASE-KR as shown:

Table 125–2—Nomenclature and clause correlation

Nomenclature	Clause and annex ^a												
	28	46		73	78	126	126	127	128	128A	129	130	130A
	Auto-Negotiation	RS	XGMII	Auto-Negotiation	EEE	2.5GBASE-T PCS/PMA	5GBASE-T PCS/PMA	2.5GBASE-X PCS/PMA	2.5GBASE-KX PMD	2.5GSEI	5GBASE-R PCS/PMA	5GBASE-KR PMD	5GSEI
2.5GBASE-KX		M	O	M	O			M	M	O			
2.5GBASE-T	M	M	O		O	M							
5GBASE-KR		M	O	M	O						M	M	O
5GBASE-T	M	M	O		O		M						

^a O = Optional, M = Mandatory.

125.2 Summary of 2.5 Gigabit and 5 Gigabit Ethernet sublayers

125.2.2 Physical coding sublayer (PCS)

Insert the following new paragraph at the end of 125.2.2:

2.5GBASE-X uses the PCS specified in Clause 127, and 5GBASE-R uses the PCS specified in Clause 129.

125.2.3 Physical Medium Attachment sublayer (PMA)

Insert following new paragraph at the end of 125.2.3:

2.5GBASE-X uses the PMA specified in Clause 127, and 5GBASE-R uses the PMA specified in Clause 129.

Change the title for 125.2.4, and insert a new subclause number (125.2.4.1) immediately afterward as shown (note that the text of the former 125.2.4 is now in 125.2.4.1):

125.2.4 Auto-Negotiation, type BASE-T

125.2.4.1 Auto-Negotiation, type BASE-T

Insert the following new subclause (125.2.4.2) after 125.2.4.1:

125.2.4.2 Auto-Negotiation, type Backplane

Auto-Negotiation (Clause 73) is used by 2.5GBASE-X and 5GBASE-R 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 differential Manchester encoding.

125.3 Delay Constraints

Change Table 125-3 as shown:

Table 125–3—Sublayer delay constraints

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
2.5GBASE-KX PHY	1 024	2	409.6	See 127.5 and 128.4
2.5GBASE-T PHY	12 800	25	5 120	Does not include delay of cable medium. See 126.11
5GBASE-KR PMD	1 024	1	102.4	See 130.4
5GBASE-R PCS PMA	3 584	7	716.8	See 129.5
5GBASE-T PHY	14 336	28	2 867.2	Does not include delay of cable medium. See 126.11

^a For 2.5GBASE-T and 2.5GBASE-X, 1 bit time (BT) is equal to 400 ps and for 5GBASE-T and 5GBASE-R, 1 bit time (BT) is equal to 200 ps. (See 1.4.160 for the definition of bit time.)

^b For 2.5GBASE-T and 2.5GBASE-X, 1 pause quantum is equal to 204.8 ns and for 5GBASE-T and 5GBASE-R, 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 the following new clauses (Clause 127 to Clause 130) in numeric order (see later in this amendment for the addition of corresponding annexes):

127. Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X

127.1 Overview

127.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) and the Physical Medium Attachment (PMA) sublayer that are common to a family of 2.5 Gb/s Physical Layer implementations known as 2.5GBASE-X. The 2.5GBASE-X PCS and 2.5GBASE-X PMA are sublayers of the 2.5 Gb/s BASE-X PHY listed in Table 125–1. The term 2.5GBASE-X is used when referring generally to Physical Layers using the PCS and PMA defined in this clause.

2.5GBASE-X PCS and PMA sublayers map the interface characteristics of the PMD sublayer (including MDI) to the services expected by the Reconciliation sublayer. 2.5GBASE-X can be extended to support any other full duplex medium requiring only that the medium be compliant at the PMD level.

127.1.2 Relationship of 2.5GBASE-X to other standards

Figure 125–1 depicts the relationships among the 2.5GBASE-X sublayers, Ethernet MAC and reconciliation layers, and the higher layers. The 2.5GBASE-X service interface is the XGMII, which is defined in Clause 46.

127.1.3 Summary of 2.5GBASE-X sublayers

Figure 127–1 shows the relationship of the 2.5GBASE-X PCS sublayer with other sublayers to the ISO Open System Interconnection (OSI) reference model.

127.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the XGMII that provides a uniform interface to the Reconciliation sublayer for all 2.5 Gb/s PHY implementations. The 2.5GBASE-X PCS provides all services required by the XGMII including encoding (decoding) of XGMII data octets to (from) ten-bit code-groups (8B/10B) for communication with the underlying PMA.

127.1.3.2 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of serial-bit-oriented physical media. The 2.5GBASE-X PMA performs the following functions:

- a) Mapping of transmit and receive code-groups between the PCS and PMA via the PMA Service Interface.
- b) Serialization (deserialization) of code-groups for transmission (reception) on the underlying serial PMD.
- c) Recovery of clock from the 8B/10B-coded data supplied by the PMD.
- d) Mapping of transmit and receive bits between the PMA and PMD via the PMD Service Interface.

IEEE Std 802.3cb-2018
 IEEE Standard for Ethernet—Amendment 1: Physical Layer Specifications and Management Parameters
 for 2.5 Gb/s and 5 Gb/s Operation over Backplane

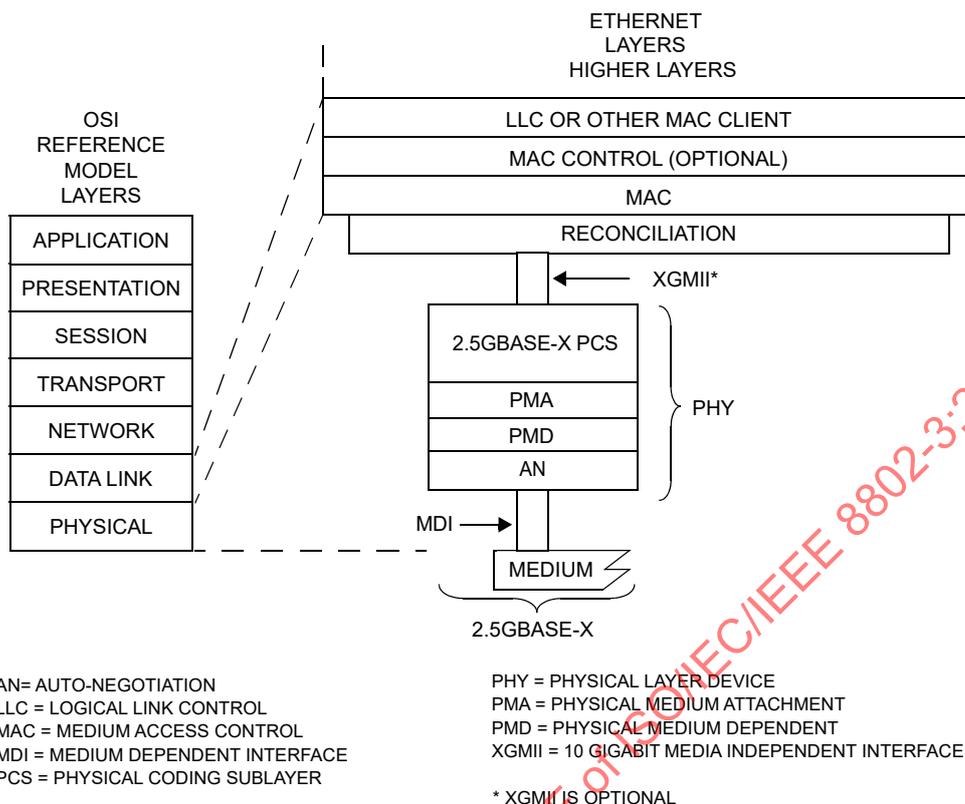


Figure 127–1—2.5GBASE-X PCS and PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

127.1.3.3 Physical Medium Attachment (PMA) service interface rates

2.5GBASE-X Physical Layer specification has nominal rate at the PMA service interface of 3.125 Gb/s, which provides MAC data rate of 2.5 Gb/s.

127.1.4 Inter-sublayer interfaces

There are a number of interfaces employed by 2.5GBASE-X. Some (such as the PMA Service Interface) use an abstract service model to define the operation of the interface. An optional physical instantiation of the PCS Interface is the XGMII. Figure 127–2 depicts the relationship and mapping of the services provided by all of the interfaces relevant to 2.5GBASE-X.

While this specification defines interfaces in terms of bits, octets, and code-groups, implementers may choose other data path widths for implementation convenience. The only exception is the XGMII, which, when implemented at an observable interconnection port, uses a 32-bit-wide data path as specified in Clause 46.

127.1.5 Functional block diagram

Figure 127–2 provides a functional block diagram of the 2.5GBASE-X PHY.

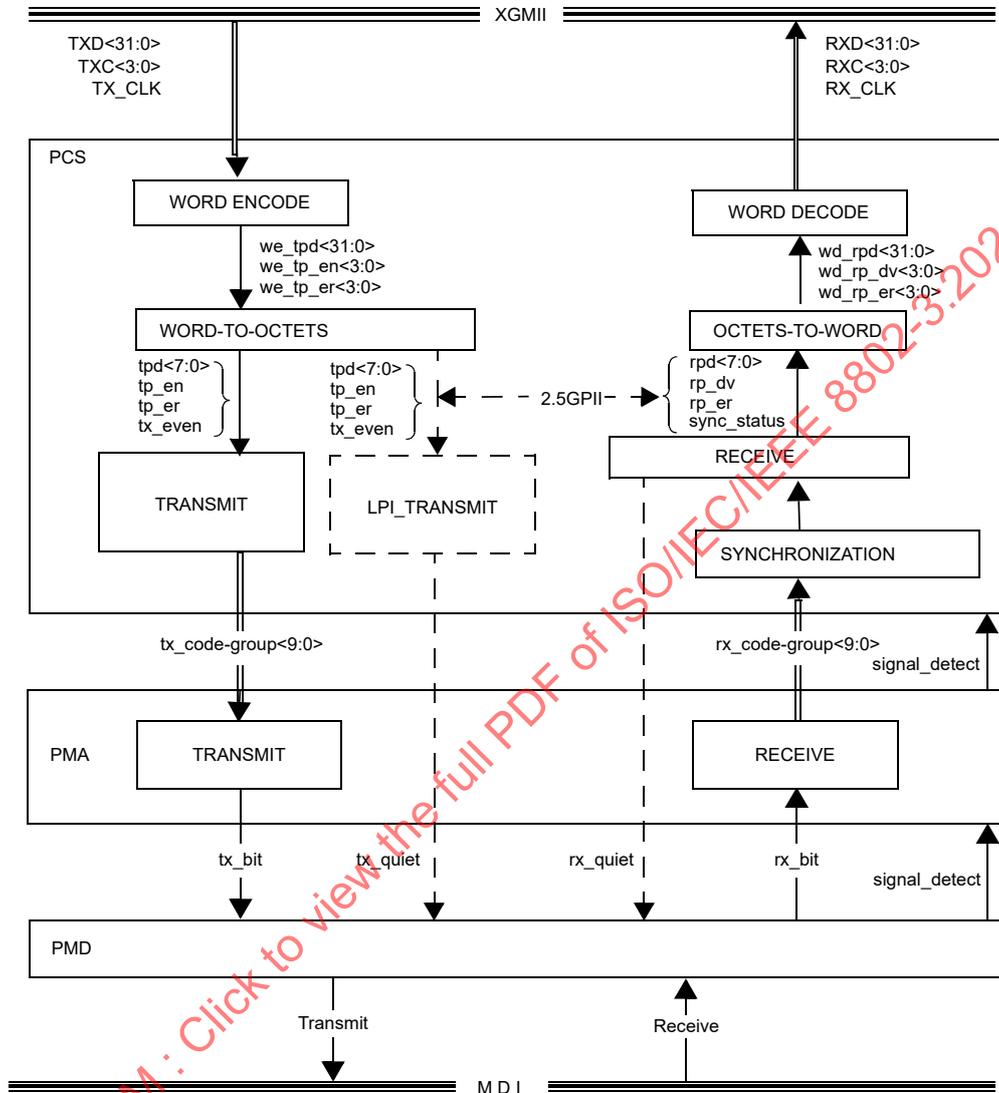


Figure 127–2—Functional block diagram of 2.5GBASE-X PHY

127.2 Physical Coding Sublayer (PCS)

127.2.1 PCS Interface (XGMII)

The PCS Service Interface allows the 2.5GBASE-X PCS to transfer information to and from a PCS client. PCS clients include the MAC (via the Reconciliation sublayer). The PCS Interface is precisely defined as the 10 Gigabit Media Independent Interface (XGMII) in Clause 46.

127.2.2 Functions within the PCS

The PCS includes the Word Encode, Word-to-Octets, Transmit, Synchronization, Receive, Octets-to-Word, and Word Decode processes for 2.5GBASE-X. The PCS shields the RS (and MAC) from the specific nature of the underlying channel.

When communicating with the XGMII, the PCS uses, in each direction, 32 data signals (TXD <31:0> and RXD <31:0>), four control signals (TXC <3:0> and RXC <3:0>), and a clock (TX_CLK and RX_CLK).

When communicating with the PMA, the PCS uses a ten-bit-wide synchronous data path, tx_code-group<9:0> and rx_code-group<9:0>, which conveys ten-bit code-groups. At the PMA Service Interface, code-group alignment and MAC packet delimiting are made possible by embedding special non-data code-groups in the transmitted code-group stream. The PCS provides the functions necessary to map packets between the XGMII format and the PMA service interface format.

The Word Encode process continuously generates four 2.5GPPII symbols (we_tpd<31:0>) and associated four bits of transmit enable (we_tp_en<3:0>) and four bits of transmit error (we_tp_er<3:0>) based upon the TXD <31:0> and TXC <3:0> signals on the XGMII, sending them to the Word-to-Octets process.

The Word-to-Octets process takes the four 2.5GPPII symbols, and associated transmit enable and transmit error, and transmits one 2.5GPPII symbol (tpd<7:0>) and its associated transmit enable (tp_en) and transmit error (tp_er) at a time to the PCS Transmit Process.

The Word-to-Octets process takes the output of the Word Encoder and presents one symbol at a time (tp_en, tp_er, tpd<7:0>) to the PCS transmit process. we_tpd<7:0> is presented first and we_tpd<31:24> is presented last.

The PCS Synchronization process continuously accepts code-groups via the PMA_UNITDATA.indication primitive and conveys received code-groups to the PCS Receive process via the SYNC_UNITDATA.indication primitive. The PCS Synchronization process sets the sync_status flag to indicate whether the PMA is functioning dependably.

The PCS Receive process continuously accepts code-groups via the SYNC_UNITDATA.indication primitive. The PCS Receive process monitors these code-groups and generates rpd<7:0>, rp_dv, and rp_er on the 2.5GPPII.

The Octets-to-Word process queues the received 2.5GPPII symbols and aligns them in groups of four 2.5GPPII symbols. Symbols may be deleted or idle symbols added in order to do the alignment.

The Word Decode process continuously accepts the four 2.5GPPII symbols from the Octets-to-Word process and generates RXD <31:0> and RXC <3:0> on the XGMII.

All PCS processes are described in detail in the state diagrams in 127.2.7.2.

127.2.3 PCS used with 2.5GBASE-KX PMD

The following requirements apply to a PCS used with a 2.5GBASE-KX PMD. Support for the Auto-Negotiation process defined in Clause 73 is mandatory. The PCS shall support the primitive AN_LINK.indication(link_status) (see 73.9). The parameter link_status shall take the value FAIL when code_sync_status=FAIL and the value OK when code_sync_status=OK. The primitive shall be generated when the value of link_status changes.

127.2.4 Use of code-groups

The transmission code used by the PCS, referred to as 8B/10B, is identical to that specified in Clause 36. The PCS maps XGMII characters into an intermediate 2.5GPll which is then mapped to 10-bit code-groups, and vice versa, using the 8B/10B block coding scheme. Implicit in the definition of a code-group is an establishment of code-group boundaries by a PCS Synchronization process. The 8B/10B transmission code as well as the rules by which the PCS ENCODE and DECODE functions generate, manipulate, and interpret code-groups are specified in 127.2.6. The XGMII to 2.5GPll mapping and vice versa are specified in 127.2.5. Code-groups and the 2.5GPll are unobservable and have no meaning outside the PCS.

127.2.5 XGMII to 2.5GPll mapping

The Word Encode, Word-to-Octets, Octets-to-Word, and Word Decode processes together define the XGMII to 2.5GPll mapping. This mapping function formats and serializes/deserializes the data between the four XGMII lanes and the single lane of the PCS transmit/receive process.

127.2.5.1 2.5 Gb/s PCS Internal Interface (2.5GPll)

The 2.5 Gb/s PCS Internal Interface (2.5GPll) is a logical interface that is internal to the 2.5GBASE-X PCS and exists solely for the purposes of defining the 2.5GBASE-X PCS functionality. Physical implementation of the 2.5GPll is optional and is not exposed outside of the PCS.

A 2.5GPll symbol is defined to be a set of *tp_en*, *tp_er*, *tpd<7:0>* variables on the transmit path, and *rp_dv*, *rp_er*, *rp_d<7:0>* variables on the receive path. The permissible encodings are shown in Table 127–1 and Table 127–2. The Word Encode and Word Decode processes maps between the four XGMII lanes to four 2.5GPll symbols. The 2.5GPll symbols are indexed as *we_tp_en<3:0>*, *we_tp_er<3:0>*, *we_tpd<31:0>*, *wd_rp_dv<3:0>*, *wd_rp_er<3:0>*, and *wd_rpd<31:0>*. The nominal rate of operation is 78.125 Msymbols/s ± 100 ppm.

The Word-to-Octets and Octets-to-Word processes serializes/de-serializes four 2.5GPll symbols, and their associated enable and error bits, to/from four consecutive single 2.5GPll symbols. The nominal rate of operation of 2.5GPll symbols is 312.5 Msymbols/s ± 100 ppm.

Table 127–1—Permissible encodings of *tpd<7:0>*, *tp_en*, *tp_er* at 2.5GPll

<i>tp_en</i>	<i>tp_er</i>	<i>tpd<7:0></i>	Description	Mnemonic
0	0	0x00 to 0xFF	Normal Inter-frame	Idle
0	1	0x00	Reserved	
0	1	0x01	Assert LPI	LPI
0	1	0x02 to 0x9B	Reserved	
0	1	0x9C	Sequence	Seq
0	1	0x9D to 0xFF	Reserved	
1	0	0x00 to 0xFF	Data	Data X
1	1	0x00 to 0xFF	Transmit Error	Err

Table 127–2—Permissible encodings of rpd<7:0>, rp_dv, rp_er at 2.5GPll

rp_dv	rp_er	rpd<7:0>	Description	Mnemonic
0	0	0x00 to 0xFF	Normal Inter-frame	Idle
0	1	0x00	Reserved	
0	1	0x01	Assert LPI	LPI
0	1	0x02 to 0x0D	Reserved	
0	1	0x0E	False carrier indication	FCI
0	1	0x0F	Carrier Extend (odd byte packets)	CE
0	1	0x10 to 0x9B	Reserved	
0	1	0x9C	Sequence	Seq
0	1	0x9D to 0xFF	Reserved	
1	0	0x00 to 0xFF	Data	Data X
1	1	0x00 to 0xFF	Receive Error	Err

127.2.5.2 Word Encode

The Word Encode process maps the four XGMII lanes (see Table 46-2) onto four 2.5GPll symbols, and their associated transmit enable and transmit error bits, as shown in Table 127–3. The XGMII encoding is specified in Table 46-3. The 2.5GPll encoding is specified in Table 127–1. The mapping of the sequence ordered set is dependent on the current state of the wencode_state variable as shown in column 5. The state of wencode_state is updated once the mapping occurs per the last column.

A sequence ordered set |Q| appears to the PMA as /K28.5/W/K28.5/W/K28.5/W/K28.5/W/. On the 2.5GPll this appears as Seq, Data S0, Seq, Data S1, Seq, Data S2, Seq, Data S3. The same sequence ordered set is assumed to be repeating over multiple XGMII cycles. Every alternating consecutive sequence ordered set on the XGMII is ignored. If a sequence ordered set occurs over odd number of cycles on the XGMII, then the final one will be truncated as Seq, Data S0, Seq, Data S1 and Seq, Data S2, Seq, Data S3 are not sent.

The 24-bit Data X, Data Y, and Data Z from the sequence ordered set is mapped to Data S0, Data S1, Data S2, Data S3 as shown in Equation (127-1).

$$\begin{aligned}
 S0<7> &= S3<7> = 0, S1<7> = S2<7> = 1 \\
 S0<5:0> &= \text{Data X}<5:0> \\
 S1<5:0> &= \text{Data Y}<3:0>, \text{Data X}<7:6> \\
 S2<5:0> &= \text{Data Z}<1:0>, \text{Data Y}<7:4> \\
 S3<5:0> &= \text{Data Z}<7:2> \\
 S_n<6> &= S_n<7> \text{ if } S_n<2> = 0 \\
 S_n<6> &= S_n<5> \text{ if } S_n<2> = 1
 \end{aligned}
 \tag{127-1}$$

The signal ordered set |Fsig| uses the same equation except S2<7> is set to 0.

Since sequence ordered set can be sent back to back it is necessary to determine the boundaries of the ordered set. {S0<7>, S1<7>, S2<7>, S3<7>} can be used to determine the boundary. {S0<7>, S1<7>} will always be 01, while the 01 combination will never occur over {S1<7>, S2<7>}, {S2<7>, S3<7>}, or {S3<7>, S0<7>}.

Only 128 combinations of S_n<7:0> are possible. When encoded to their 10-bit equivalent, these 128 code-groups are defined to be in the set of /W/ (see 127.2.7.1.2).

Table 127-3—Word Encode mapping

XGMII				2.5GPII				wencode _state (n+1)	
Lane 0	Lane 1	Lane 2	Lane 3	wencode _state (n)		wd_tpd<7:0> we_tp_en<0> we_tp_er<0>	wd_tpd<15:8> we_tp_en<1> we_tp_er<1>	wd_tpd<23:16> we_tp_en<2> we_tp_er<2>	wd_tpd<31:24> we_tp_en<3> we_tp_er<3>
Data A or Err	Data B or Err	Data C or Err	Data D or Err	Don't care	⇒	Data A or Err	Data B or Err	Data C or Err	Data D or Err
Idle	Idle	Idle	Idle	Don't care	⇒	Idle	Idle	Idle	Idle
LPI	LPI	LPI	LPI	Don't care	⇒	LPI	LPI	LPI	LPI
Start	Data A or Err	Data B or Err	Data C or Err	Don't care	⇒	Data 0x55	Data A or Err	Data B or Err	Data C or Err
Terminate	Idle	Idle	Idle	Don't care	⇒	Idle	Idle	Idle	Idle
Data A or Err	Terminate	Idle	Idle	Don't care	⇒	Data A or Err	Idle	Idle	Idle
Data A or Err	Data B or Err	Terminate	Idle	Don't care	⇒	Data A or Err	Data B or Err	Idle	Idle
Data A or Err	Data B or Err	Data C or Err	Terminate	Don't care	⇒	Data A or Err	Data B or Err	Data C or Err	Idle
Sequence	Data X	Data Y	Data Z	IDLE	⇒	Seq	Data S0	Seq	Data S1
Sequence	Data X	Data Y	Data Z	SEQ	⇒	Seq	Previous Data S2 ^a	Seq	Previous Data S3 ^a
Sequence	Data X	Data Y	Data Z	DATA	⇒	Idle	Idle	Idle	Idle
else				Don't care	⇒	Err	Err	Err	Err

^a Previous Data S2 and Previous Data S3 are the values of S2 and S3 respectively as calculated by Equation (127-1) during the previous mapping.

127.2.5.3 Word-to-Octets

The Word-to-Octets process takes the output of the Word Encoder and presents one symbol at a time (tp_en , tp_er , $tpd\langle 7:0 \rangle$) to the PCS transmit process. $we_tpd\langle 7:0 \rangle$ is presented first and $we_tpd\langle 31:24 \rangle$ is presented last.

The Word-to-Octets process shall be synchronized to the PCS transmit process such that the $we_tpd\langle 7:0 \rangle$ and $we_tpd\langle 23:16 \rangle$ symbols are presented to the PCS transmit process which will result in the corresponding ordered set to be output to the PMA when the variable tx_even is TRUE and $we_tpd\langle 15:8 \rangle$ and $we_tpd\langle 31:24 \rangle$ variables when tx_even is FALSE.

127.2.5.4 Octets-to-Word

The Octets-to-Word process de-serializes the output of the PCS receive process to four symbols ($wd_rp_dv\langle 3:0 \rangle$, $wd_rp_er\langle 3:0 \rangle$, $wd_rpd\langle 31:0 \rangle$). $wd_rpd\langle 7:0 \rangle$ is the earliest to arrive and $wd_rpd\langle 31:24 \rangle$ is the last.

The Octets-to-Word process inserts idle symbols or deletes symbols from the sequence of $rpd\langle 7:0 \rangle$ symbols received to achieve the following conditions:

- A transition of a 2.5GPPII idle symbol to a data or error symbol shall place the data or error symbol on $wd_rpd\langle 7:0 \rangle$.
- A transition of a 2.5GPPII idle symbol to a LPI symbol shall place the LPI symbol on either $wd_rpd\langle 7:0 \rangle$ or $wd_rpd\langle 23:16 \rangle$.
- A transition of a 2.5GPPII LPI symbol to an idle symbol shall place the idle symbol on either $wd_rpd\langle 7:0 \rangle$ or $wd_rpd\langle 23:16 \rangle$.
- The start of $\|Q\|$ or $\|Fsig\|$ set shall always occur on $wd_rpd\langle 7:0 \rangle$. 127.2.5.2 describes how the start of $\|Q\|$ and $\|Fsig\|$ can be determined.

The Octets-to-Word process maintains a Deficit Idle Count (DIC) that represent the cumulative count of idle symbols in $rpd\langle 7:0 \rangle$ added or deleted from the sequence of received symbols. The DIC is incremented by one for each idle symbol deleted and decremented by one for each idle symbol added. The DIC shall be bounded to a minimum of 0 and a maximum of 3.

Note that in a properly behaved system, deletion of idle symbols from $rpd\langle 7:0 \rangle$ onto $wd_rpd\langle 31:0 \rangle$ should only occur at most once at the beginning of link, and afterwards no further insertions or deletions are required. In order to interoperate with the application described in Annex 127A, additional symbol insertions and deletions may be required during normal operation.

The only symbol that may be inserted is a idle symbol. However any symbol may be deleted. Usually this will either be idle or LPI symbols, though in pathological error conditions (e.g., unterminated packet followed immediately with sequence ordered set) some other symbol may be deleted.

127.2.5.5 Word Decode

The Word Decode process maps the four 2.5GPll symbols onto the four XGMll lanes (see Table 46-2) as shown in Table 127-4. The XGMll encoding is specified in Table 46-4. The seq_s2_s3 variable indicates whether the next four 2.5GPll symbols are of the form Sequence, Data, Sequence, Data. The 2.5GPll encoding is specified in Table 127-2. The mapping is dependent on the current state of the wdecode_state and next_seq_s2_s3 variables as shown in columns 5 and 6. The state of wdecode_state is updated once the mapping occurs per the last column.

The 24-bit Data X, Data Y, and Data Z from the sequence ordered is reconstructed from Data S0, Data S1, Data S2, Data S3 according to Equation (127-2).

$$\begin{aligned} \text{Data X}<7:0> &= \text{S1}<1:0>, \text{S0}<5:0> \\ \text{Data Y}<7:0> &= \text{S2}<3:0>, \text{S1}<5:2> \\ \text{Data Z}<7:0> &= \text{S3}<5:0>, \text{S2}<5:4> \end{aligned} \tag{127-2}$$

127.2.6 8B/10B transmission code

The transmission code used by the PCS, referred to as 8B/10B, is identical to that specified in Clause 36. In addition to the requirements in this clause, a 2.5GBASE-X PCS shall also meet the 8B/10B transmission code requirements specified in 36.2.4.1 through 36.2.4.6, 36.2.4.8, and 36.2.4.9. The relationship of code-group bit positions to PMA and other PCS constructs is illustrated in Figure 127-3.

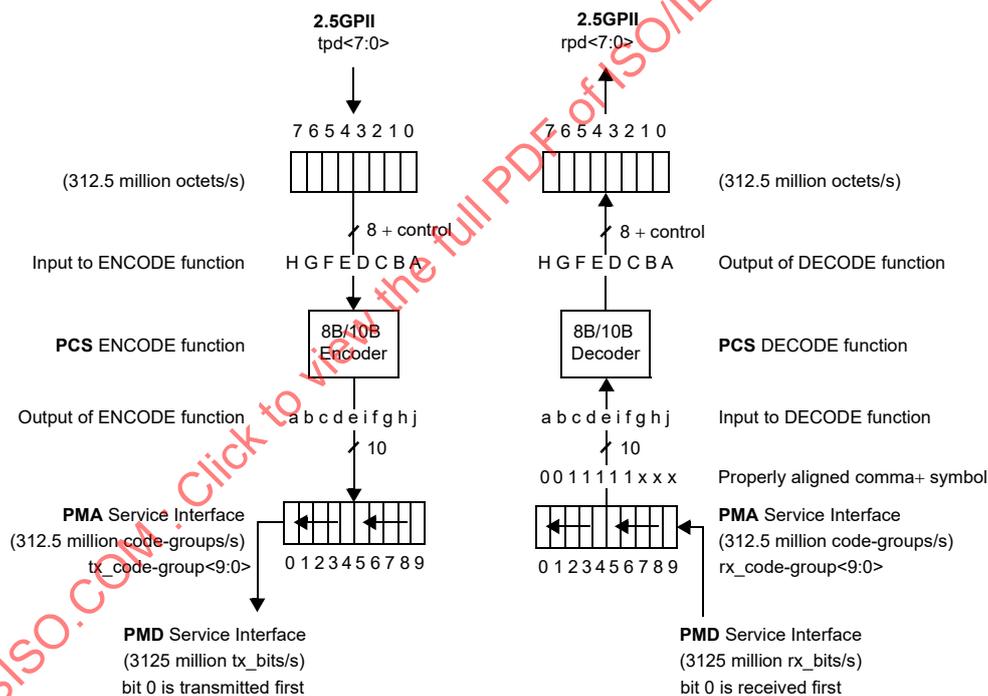


Figure 127-3—PCS 8B/10B reference diagram

127.2.6.1 Notation conventions

The 8B/10B transmission code uses letter notation for describing the bits of an unencoded information octet and a single control variable according to 36.2.4.1.

127.2.6.2 Transmission order

Code-group bit transmission order is illustrated in Figure 127-3 and defined in 36.2.4.2.

Table 127-4—Word Decode mapping

2.5GPII				wdecode _state (n)	next seq_s2_s3 ^a	XGMII				wdecode _state (n+1)	
wd_rpd<7:0> we_rp_en<0> we_rp_er<0>	wd_rpd<15:8> we_rp_en<1> we_rp_er<1>	wd_rpd<23:16> we_rp_en<2> we_rp_er<2>	wd_rpd<31:24> we_rp_en<3> we_rp_er<3>			Lane 0	Lane 1	Lane 2	Lane 3		
Data A or Err	Data B or Err	Data C or Err	Data D or Err	≠ IDLE	Don't care	⇒	Data A or Err	Data B or Err	Data C or Err	Data D or Err	DATA
Data	Data A or Err	Data B or Err	Data C or Err	IDLE	Don't care	⇒	SOP	Data A or Err	Data B or Err	Data C or Err	DATA
Idle	Idle	Idle	Idle	≠ DATA	Don't care	⇒	Idle	Idle	Idle	Idle	IDLE
Idle	Idle	Idle	Idle	DATA	Don't care	⇒	Terminate	Idle	Idle	Idle	IDLE
Data A or Err	Idle or CE	Idle	Idle	DATA	Don't care	⇒	Data A or Err	Terminate	Idle	Idle	IDLE
Data A or Err	Data B or Err	Idle	Idle	DATA	Don't care	⇒	Data A or Err	Data B or Err	Terminate	Idle	IDLE
Data A or Err	Data B or Err	Data C or Err	Idle or CE	DATA	Don't care	⇒	Data A or Err	Data B or Err	Data C or Err	Terminate	IDLE
LPI	LPI	LPI	LPI	≠ DATA	Don't care	⇒	LPI	LPI	LPI	LPI	IDLE
Idle	Idle	LPI	LPI	Don't care	Don't care	⇒	LPI	LPI	LPI	LPI	IDLE
LPI	LPI	Idle	Idle	≠ DATA	Don't care	⇒	Idle	Idle	Idle	Idle	IDLE
Seq	Data S0	Seq	Data S1	≠ DATA	TRUE	⇒	Sequence	Data X	Data Y	Data Z	SEQ
Seq	Data S0	Seq	Data S1	≠ DATA	FALSE	⇒	Idle	Idle	Idle	Idle	IDLE
Seq	Data S2	Seq	Data S3	SEQ	Don't care	⇒	Sequence	Data X	Data Y	Data Z	IDLE
else				Don't care	Don't care	⇒	Err	Err	Err	Err	ERR

^a next_seq_s2_s3 is TRUE when the next four GPII octets represent the S2 and S3 sequence ordered set and FALSE otherwise.

127.2.6.3 Generating code-groups and checking the validity of received code

Valid code-groups are defined in 36.2.4.3.
 The running disparity rules are defined in 36.2.4.4.
 The code-group generation is defined in 36.2.4.5.
 The check on the validity of received code-groups is defined in 36.2.4.6.

127.2.6.4 Ordered sets

Table 127–5 lists the defined ordered sets, consisting of a single special code-group or combinations of special and data code-groups. Ordered sets which include /K28.5/ provide the ability to obtain bit and code-group synchronization and establish ordered set alignment (see 36.2.4.9 and 127.3.2.4). Ordered sets provide for the delineation of a packet and synchronization between the transmitter and receiver circuits at opposite ends of a link. Certain PHYs include an option (see 78.3) to transmit or receive /LI/, /LI1/ and /LI2/ to support Energy-Efficient Ethernet (see Clause 78).

Table 127–5—Defined ordered sets

Code	Ordered Set	Number of Code-Groups	Encoding
/I/	IDLE		Correcting /I1/, Preserving /I2/
/I1/	IDLE 1	2	/K28.5/D5.6/
/I2/	IDLE 2	2	/K28.5/D16.2/
	Encapsulation		
/R/	Carrier_Extend	1	/K23.7/
/S/	Start_of_Packet	1	/K27.7/
/T/	End_of_Packet	1	/K29.7/
/V/	Error_Propagation	1	/K30.7/
/LI/	LPI		Correcting /LI1/, Preserving /LI2/
/LI1/	LPI 1	2	/K28.5/D6.5/
/LI2/	LPI 2	2	/K28.5/D26.4/
	Link Status		
/Q/	Sequence ordered set	8	/K28.5/W0/K28.5/W1/K28.5/W2/ K28.5/W3 ^a
	Reserved		
/Fsig/	Signal ordered set	8	/K28.5/W0/K28.5/W1/K28.5/W2/ K28.5/W3 ^a

^a /W0/, /W1/, /W2/, /W3/ are the 10-bit /Dx.y/ version of S0, S1, S2, S3 as defined per 127.2.5.2 and will never have a value of /D5.6/, /D16.2/, /D6.5/, /D26.4/.

Ordered sets are specified according to the following rules:

- a) Ordered sets consist of either one, two, or eight code-groups.
- b) The first code-group of all ordered sets is always a special code-group.
- c) The second code-group of all multi-code-group ordered sets is always a data code-group. The second code-group is used to distinguish the ordered set from all other ordered sets.

127.2.6.5 Comma considerations

The comma considerations are described in 36.2.4.8 and 36.2.4.9.

127.2.6.6 Sequence (/Q/)

A sequence ordered set is used to convey various optional link status such as local fault or remote fault. The 24-bit data of the sequence ordered set on the XGMII, when implemented, are mapped to S0, S1, S2, S3 (see 127.2.5.2), and /W0/, /W1/, /W2/, /W3/ are the 8B/10B mapped version.

S0, S1, S2, S3 are constructed such that each /Wn/ can be mapped to only 128 out of 256 possible /Dx.y/ code-groups. /W/ is defined to be the set of 128 /Dx.y/ code-groups that can appear in a sequence ordered set. /W/ does not contain /D5.6/, /D16.2/, /D6.5/, /D26.4/.

It is possible for the transmitter to send out a truncated sequence ordered set that appears as /K28.5/W0/K28.5/W1/. When a truncated sequence ordered set is received, the Word Decode process will convert it to idles.

127.2.6.7 Data (/D/)

A data code-group, when not used to distinguish or convey information for a defined ordered set, conveys one octet of arbitrary data supplied on the XGMII. The sequence of data code-groups is arbitrary, where any data code-group can be followed by any other data code-group. Data code-groups are coded and decoded but not interpreted by the PCS. Successful decoding of the data code-groups depends on proper receipt of the Start_of_Packet delimiter, as defined in 127.2.6.10 and the checking of validity, as defined in 36.2.4.6.

127.2.6.8 IDLE (/I/)

IDLE ordered sets (/I/) are transmitted continuously and repetitively whenever the XGMII is idle. /I/ provides a continuous fill pattern to establish and maintain clock synchronization. /I/ is emitted from, and interpreted by, the PCS. /I/ consists of one or more consecutively transmitted /I1/ or /I2/ ordered sets, as defined in Table 127-5.

The /I1/ ordered set is defined such that the running disparity at the end of the transmitted /I1/ is opposite that of the beginning running disparity. The /I2/ ordered set is defined such that the running disparity at the end of the transmitted /I2/ is the same as the beginning running disparity. The first /I/ following a packet or a sequence ordered set restores the current positive or negative running disparity to a negative value. All subsequent /I/s are /I2/ to ensure negative ending running disparity.

Distinct carrier events are separated by /I/s.

A received ordered set that consists of two code-groups, the first of which is /K28.5/ and the second of which is a data code-group other than any of the 128 possible /W/, /D21.5/, or /D2.2/ (or /D6.5/ or /D26.4/ to support IEEE capability), is treated as an /I/ ordered set.

127.2.6.9 Low Power Idle (/LPI/)

LPI is transmitted in the same manner as IDLE. LPI ordered sets (/LPI/) are transmitted continuously and repetitively whenever the XGMII is indicating “Assert LPI”.

127.2.6.10 Start_of_Packet delimiter (SPD)

A Start_of_Packet delimiter (SPD) is used to delineate the starting boundary of a data transmission sequence. Upon initiation of packet transmission, the PCS replaces the first octet of the MAC preamble with

SPD. Upon initiation of packet reception, the PCS replaces the received SPD delimiter with the data octet value associated with the first preamble octet. A SPD delimiter consists of the code-group /S/, as defined in Table 127–5.

127.2.6.11 End_of_Packet delimiter (EPD)

An End_of_Packet delimiter (EPD) is used to delineate the ending boundary of a packet. The EPD is transmitted by the PCS following the last data octet comprising the FCS of the MAC packet. On reception, EPD is interpreted by the PCS as terminating a packet. A EPD consists of the code-groups /T/R/. If /R/ is transmitted in an even-numbered code-group position, the PCS appends a single additional /R/ to the code-group stream to ensure that the subsequent /I/ is aligned on an even numbered code-group boundary. An /I/ always follows the conclusion of EPD. The /T/R/ or /T/R/R/ occupies part of the region considered by the MAC to be the IPG. The code-group /T/ and /R/ are defined in Table 127–5.

127.2.6.12 Error_Propagation (/V/)

Error_Propagation (/V/) indicates that the PCS client wishes to indicate a transmission error to its peer entity. /V/ is emitted from the PCS when the XGMII indicates transmit error propagation, or when the XGMII presents a mapping that is undefined in Table 127–3 to the Word Encode process. The code group /V/ is defined in Table 127–5.

The presence of Error_Propagation or any invalid code-group on the medium denotes an error condition. Invalid code-groups are not intentionally transmitted onto the medium.

127.2.7 Detailed functions and state diagrams

The body of this subclause comprises 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.

127.2.7.1 State variables

127.2.7.1.1 Notation conventions

/x/ Denotes the constant code-group specified in 127.2.7.1.2 (valid code-groups must follow the rules of running disparity as per 127.2.6.3).

[/x/] Denotes the latched received value of the constant code-group (/x/) specified in 127.2.7.1.2 and conveyed by the SYNC_UNITDATA.indication message described in 127.2.7.1.6.

127.2.7.1.2 Constants

/COMMA/ The set of special code-groups which include a comma as specified in 36.2.4.9 and listed in Table 36-2.

/D/ The set of 256 code-groups corresponding to valid data, as specified in 127.2.6.7.

/Dx.y/ One of the set of 256 code-groups corresponding to valid data, as specified in 127.2.6.7.

- /I/**
The IDLE ordered set group, comprising either the /I1/ or /I2/ ordered sets, as specified in 127.2.6.8.
- /INVALID/**
The set of invalid data or special code-groups, as specified in 36.2.4.6.
- /Kx.y/**
One of the set of 12 code-groups corresponding to valid special code-groups, as specified in Table 36-2.
- /Q/**
Sequence ordered set as specified in 127.2.6.6. A properly formed sequence order set appears as /K28.5/W/K28.5/W/K28.5/W/K28.5/W/. A truncated sequence order set appears as /K28.5/W/K28.5/W/.
- /R/**
The code-group used for the second and, if present, the third code-group in an End_of_Packet delimiter as specified in 127.2.6.11.
- /S/**
The code-group corresponding to the Start_of_Packet delimiter (SPD) as specified in 127.2.6.10.
- /T/**
The code-group used for the first code-group in the End_of_Packet delimiter as specified in 127.2.6.11.
- /V/**
The Error_Propagation code-group, as specified in 127.2.6.12.
- /W/**
The set of 128 code-groups that is generated by ENCODE(s<7:0>) where for all 128 possible values of x<6:0>
 $s<7> = x<6>$, $s<5:0> = x<5:0>$;
 $s<6>$ is set to $x<5>$ when $x<2> = 1$; and $s<6>$ is set to $x<6>$ when $x<2> = 0$
 NOTE— /W/ is a subset of /D/.

PL_LIMIT

The number of 2.5GPII symbols to preload in the Octets-to-Word process after release from the RESET state. The number of symbols to preload is implementation dependent and should be minimized to reduce latency. The minimum number must be 3 to account for the deficit idle counting in 127.2.5.4.

The following constant is used only for the EEE capability:

- /LI/**
The LP_IDLE ordered set group, comprising either the /LI1/ or /LI2/ ordered sets, as specified in 127.2.6.9.

127.2.7.1.3 Variables**assert_seq**

Alias used for sequence ordered set, consisting of the following terms:

$tp_en=0 * tp_er=1 * (tpd<7:0> = 0x9C)$

cgbad

Alias for the following terms: $((rx_code_group \in /INVALID/) + (rx_code_group = /COMMA/*rx_even=TRUE)) * PMA_UNITDATA.indication$

cgood

Alias for the following terms: $!((rx_code_group \in /INVALID/) + (rx_code_group = /COMMA/*rx_even=TRUE)) * PMA_UNITDATA.indication$

EVEN

The latched state of the rx_even variable, when rx_even=TRUE, as conveyed by the SYNC_UNITDATA.indication message described in 127.2.7.1.6.

mr_loopback

A Boolean that indicates the enabling and disabling of data being looped back through the PHY. Loopback of data through the PHY is enabled when Control register bit 3.0.14 is set to one.

Values: FALSE; Loopback through the PHY is disabled.
TRUE; Loopback through the PHY is enabled.

mr_main_reset

Controls the resetting of the PCS via Control Register bit 3.0.15.

Values: FALSE; Do not reset the PCS.
TRUE; Reset the PCS.

ODD

The latched state of the rx_even variable, when rx_even=FALSE, as conveyed by the SYNC_UNITDATA.indication message described in 127.2.7.1.6.

power_on

Condition that is true until such time as the power supply for the device that contains the PCS has reached the operating region. The condition is also true when the device has low power mode set via Control register bit 3.0.11.

Values: FALSE; The device is completely powered (default).
TRUE; The device has not been completely powered.

NOTE—Power_on evaluates to its default value in each state where it is not explicitly set.

rpd<7:0>

Single lane 2.5GPII receive data from the PCS receive process.

wd_rpd<31:0>

Receive data from the Octets-to-Word process. x = 0, 1, 2, 3 for the four sets of 2.5GPII.

rp_dv

Single lane 2.5GPII receive data valid from the PCS receive process.

Values: 0 or 1.

wd_rp_dv<3:0>

Receive data valid from the Octets-to-Word process.

Values: 0 or 1.

rp_er

Single lane 2.5GPII receive error from the PCS receive process.

Values: 0 or 1.

wd_rp_er<3:0>

Receive error from the Octets-to-Word process.

Values: 0 or 1.

rx_bit

A binary parameter conveyed by the PMD_UNITDATA.indication service primitive, as specified in 128.2.2, to the PMA.

Values: ZERO; Data bit is a logical zero.
ONE; Data bit is a logical one.

rx_code-group<9:0>

A 10-bit vector represented by the most recently received code-group from the PMA. The element rx_code-group<0> is the least recently received (oldest) rx_bit; rx_code-group<9> is the most recently received rx_bit (newest). When code-group alignment has been achieved, this vector contains precisely one code-group.

rx_even

A Boolean set by the PCS Synchronization process to designate received code-groups as either even- or odd-numbered code-groups as specified in 127.2.6.2.

Values: TRUE; Even-numbered code-group being received.
FALSE; Odd-numbered code-group being received.

RXC<3:0>

The RXC<3:0> signal of the XGMII as specified in Clause 46. Set by the Word Decode process.

RXD<31:0>

The RXD<31:0> signal of the XGMII as specified in Clause 46. Set by the Word Decode process.

signal_detect

A Boolean set by the PMD continuously via the PMD_SIGNAL.indication(SIGNAL_DETECT) message to indicate the status of the incoming link signal.

Values: FAIL; A signal is not present on the link.
OK; A signal is present on the link.

sync_status

Alias used by the PCS receive state diagram, consisting of the following terms:

$sync_status = code_sync_status + rx_lpi_active$.

Values: FAIL; A signal is not present on the link.
OK; A signal is present on the link.

NOTE—If IEEE 802.3EEE is not supported, the variable rx_lpi_active is always false, and this variable is identical to code_sync_status controlled by the synchronization state diagram.

tpd<7:0>

Single lane 2.5GPII transmit data to the PCS transmit process.

tpd_t1<7:0>

The value of tpd<7:0> latched by cg_timer_done = TRUE.

Values: 0 or 1.

we_tpd<31:0>

Transmit data output of the WORD ENCODE process.

- tp_en**
Single lane 2.5GPII transmit data enable to the PCS transmit process.
Values: 0 or 1.
- we_tp_en<3:0>**
Transmit data valid output of the WORD ENCODE process.
Values: 0 or 1.
- tp_er**
Single lane 2.5GPII transmit error to the PCS transmit process.
Values: 0 or 1.
- we_tp_er<3:0>**
Transmit error output of the WORD ENCODE process.
Values: 0 or 1.
- tx_bit**
A binary parameter used to convey data from the PMA to the PMD via the PMD_UNITDATA.request service primitive as specified in 128.2.1.
Values: ZERO; Data bit is a logical zero.
ONE; Data bit is a logical one.
- tx_code-group<9:0>**
A vector of bits representing one code-group, as specified in Table 36-1a through Table 36-1e, or Table 36-2, which has been prepared for transmission by the PCS Transmit process. This vector is conveyed to the PMA as the parameter of the PMD_UNITDATA.request(tx_bit) service primitive. The element tx_code-group<0> is the first tx_bit transmitted; tx_code-group<9> is the last tx_bit transmitted.
- tx_disparity**
A Boolean set by the PCS Transmit process to indicate the running disparity at the end of code-group transmission as a binary value. Running disparity is described in 36.2.4.4.
Values: POSITIVE
NEGATIVE
- tx_even**
A Boolean set by the PCS Transmit process to designate transmitted code-groups as either even or odd numbered code-groups as specified in 127.2.6.2.
Values: TRUE; Even-numbered code-group being transmitted.
FALSE; Odd-numbered code-group being transmitted.
- tx_o_set**
One of the following defined ordered sets: /T/, /R/, /I/, /S/, /V/, /LI/, or one of the following code-groups: /K28.5/ or /D/.
- TXC<3:0>**
The TXC<3:0> signal of the XGMII as specified in Clause 46.
- TXD<31:0>**
The TXD<31:0> signal of the XGMII as specified in Clause 46.

wdecode_state

Word Decoder State used by the Word Decode process to properly assemble the next XGMII value to output.

Values: DATA;
IDLE;
SEQ;

wencode_state

Word Encoder State used by the Word Encode process to properly assemble the Sequence ordered set.

Values: DATA;
IDLE;
SEQ;

xgmii_txc_lo<3:0>

The value of xgmii_txc_lo_neg<3:0> latched by the rising edge of TX_CLK.

xgmii_txc_lo_neg<3:0>

The value of TXC<3:0> latched by the falling edge of TX_CLK.

xgmii_txc_hi<3:0>

The value of TXC<3:0> latched by the rising edge of TX_CLK.

xgmii_txd_lo<31:0>

The value of xgmii_txd_lo_neg<31:0> latched by the rising edge of TX_CLK.

xgmii_txd_lo_neg<31:0>

The value of TXD<3:0> latched by the falling edge of TX_CLK.

xgmii_txd_hi<31:0>

The value of TXD<31:0> latched by the rising edge of TX_CLK.

The following variables are used only for the EEE capability:

assert_lpidle

Alias used for the optional LPI function, consisting of the following terms:

$(tp_en=0 * tp_er=1 * (tpd<7:0> = 0x01))$

code_sync_status

A parameter set by the PCS Synchronization process to reflect the status of the link as viewed by the receiver.

Values: FAIL; The receiver is not synchronized to code-group boundaries.
OK; The receiver is synchronized to code-group boundaries.

idle_d

Alias for the following terms:

SUDI(![D21.5/] * ![D2.2/]) when EEE is not supported, or

SUDI(![D21.5/] * ![D2.2/] * ![D6.5/] * ![D26.4/]) when EEE is supported.

rx_lpi_active

A Boolean variable that is set to TRUE when the receiver is in a low power state and set to FALSE when it is in an active state and capable of receiving data.

rx_quiet

A Boolean variable set to TRUE while in the RX_QUIET state and set to FALSE otherwise.

tx_quiet

A Boolean variable set to TRUE when the transmitter is in the TX_QUIET state and set to FALSE otherwise. When set to TRUE, the PMD will disable the transmitter as described in 128.6.5.

127.2.7.1.4 Functions

carrier_detect

In the PCS Receive process, this function uses for input the latched code-group ([/x/]) and latched rx_even (EVEN/ODD) parameters of the SYNC_UNITDATA.indication message from the PCS Synchronization process. When SYNC_UNITDATA.indication message indicates EVEN, the carrier_detect function detects carrier when either:

- a) A two or more bit difference between [/x/] and both /K28.5/ encodings exists (see Table 36-2); or
- b) A two to nine bit difference between [/x/] and the expected /K28.5/ (based on current running disparity) exists.

Values: TRUE; Carrier is detected.
 FALSE; Carrier is not detected.

check_end

Present End_of_Packet function used by the PCS Receive process. The check_end function returns the current and next two code-groups in rx_code-group<9:0>.

DECODE([/x/])

In the PCS Receive process, this function takes as its argument the latched value of rx_code-group<9:0> ([/x/]) and the current running disparity, and returns the corresponding 2.5GPII rpd<7:0>, as per Table 36-1a–e. DECODE also updates the current running disparity per the running disparity rules outlined in 36.2.4.4.

ENCODE(x)

In the PCS Transmit process, this function takes as its argument (x), where x is a 2.5GPII tpd<7:0> octet, and the current running disparity, and returns the corresponding ten-bit code-group as per Table 36-1a–e. ENCODE also updates the current running disparity variable tx_disparity per the running disparity rules outlined in 36.2.4.4.

NEXTSEQ

Present function used by the Word Decode process that returns whether the next four 2.5GPII symbols presented to the Word Decode process is of the form: Sequence, Data, Sequence, Data.

Values: TRUE; Next four 2.5GPII symbols are Sequence, Data, Sequence, Data.
 FALSE; Next four 2.5GPII symbols are not Sequence, Data, Sequence, Data.

signal_detectCHANGE

In the PCS Synchronization process, this function monitors the signal_detect variable for a state change. The function is set upon state change detection.

Values: TRUE; A signal_detect variable state change has been detected.
 FALSE; A signal_detect variable state change has not been detected (default).

NOTE—signal_detectCHANGE is set by this function definition; it is not set explicitly in the state diagrams. signal_detectCHANGE evaluates to its default value upon state entry.

SINSERT(x)

Add a single 2.5GPII symbol to the end of a queue that stores the 2.5GPII symbol presented by the receive process. The variable x is the 2.5GPII symbol (rp_dv, rp_er, rpd<7:0>). If x is a null set, then all content in the queue is emptied. The depth of the queue is implementation dependent.

VOID(x)

$x \in /D/, /T/, /R/, /K28.5/$. Substitutes /V/ on a per code-group basis as requested by the 2.5GPII.

If [tp_en=0 * tp_er=1 * tpd<7:0>≠(0000 1111)],

then return /V/;

Else if [tp_en=1 * tp_er=1],

then return /V/;

Else return x.

WALIGN

In the PCS Octets-to-Word process, this function performs the alignment according to 127.2.5.4. Four 2.5GPII symbols (wd_rp_dv<3:0>, wd_rp_er<3:0>, wd_rpd<31:0>) are returned by this function. wd_rpd<7:0> is the earliest to arrive and wd_rpd<31:24> is last. The SINSERT(x) function adds 2.5GPII symbols to the queue. The WALIGN function removes one to seven 2.5GPII symbols from the front of the queue every time it is called.

When no 2.5GPII symbols are inserted or deleted, this function will return the first four 2.5GPII symbols in the queue and remove them from the queue.

When X 2.5GPII idles symbols are inserted, then X 2.5GPII idles symbols and the first 4 - X 2.5GPII symbols are returned by the function, and the first 4 - X 2.5GPII symbols are removed from the queue.

When X 2.5GPII symbols are deleted, then the first X 2.5GPII symbols are removed from the queue, and the next four in the queue are returned by the function and then removed from the queue.

WDECODE(x, y, z)

In the PCS Word Decode process, this function performs the mapping according to 127.2.5.5. The variable x is four sets of 2.5GPII variables wd_rp_dv<3:0>, wd_rp_er<3:0>, wd_rpd<31:0>, the variable y is the current state of the wdecode_state variable, and the variable z indicates whether the next four 2.5GPII symbols are of the form: Sequence, Data, Sequence, Data.

The output is XGMII RXC<3:0>, RXD<31:0>, wdecode_state.

WENCODE(x, y)

In the PCS Word Encode process, this function performs the mapping according to 127.2.5.2. The variable x is xgmii_txc_lo<3:0> or xgmii_txc_hi<3:0>, xgmii_txd_lo<31:0> or xgmii_txd_hi<31:0>, and the variable y is the current state of the wencode_state variable.

The output is four sets of 2.5GPII variables and the updated state of the wencode_state variable and is ordered as follows: we_tp_en<3:0>, we_tp_er<3:0>, we_tpd<31:0>, wencode_state.

127.2.7.1.5 Counters

good_cgs

Count of consecutive valid code-groups received.

plcnt

Count of number the number of 2.5GPPII symbols to preload in the Octets-to-Word process after release from the RESET state. The number of symbols to preload is implementation dependent and should be minimized to reduce latency.

The following counter is used only for the EEE capability:

wake_error_counter

A counter that is incremented each time that the LPI receive state diagram enters the RX_WTF state indicating that a wake time fault has been detected. The counter is reflected in register 3.22 (see 45.2.3.10).

127.2.7.1.6 Messages

PMA_UNITDATA.indication(rx_code-group<9:0>)

A signal sent by the PMA Receive process conveying the next code-group received over the medium (see 127.3.1.2).

PMA_UNITDATA.request(tx_code-group<9:0>)

A signal sent to the PMA Transmit process conveying the next code-group ready for transmission over the medium (see 127.3.1.1).

PMD_SIGNAL.indication(SIGNAL_DETECT)

A signal sent by the PMD to indicate the status of the signal being received on the MDI.

PUDI

Alias for PMA_UNITDATA.indication(rx_code-group<9:0>).

PUDR

Alias for PMA_UNITDATA.request(tx_code-group<9:0>).

SUDI

Alias for SYNC_UNITDATA.indication(parameters).

SYNC_UNITDATA.indication(parameters)

A signal sent by the PCS Synchronization process to the PCS Receive process conveying the following parameters:

Parameters: [x/]; the latched value of the indicated code-group (/x/);
EVEN/ODD; the latched state of the rx_even variable;

Value: EVEN; passed when the latched state of rx_even=TRUE.
ODD; passed when the latched state of rx_even=FALSE.

TX_OSET.indication

A signal sent to the PCS Transmit ordered set process from the PCS Transmit code-group process signifying the completion of transmission of one ordered set.

The following messages are used only for the EEE capability:

PMD_RXQUIET.request(rx_quiet)

A signal sent by the PCS/PMA LPI receive state diagram to the PMD. This message is ignored by devices that do not support EEE capability.

Values: TRUE: The receiver is in a quiet state and is not expecting incoming data.
FALSE: The receiver is ready to receive data.

PMD_TXQUIET.request(tx_quiet)

A signal sent by the PCS/PMA LPI transmit state diagram to the PMD. This message is ignored by devices that do not support the optional LPI mechanism.

Values: TRUE: The transmitter is in a quiet state and may cease to transmit a signal on the medium.
 FALSE: The transmitter is ready to transmit data.

127.2.7.1.7 Timers

cg_timer

A continuous free-running timer. If XGMII is implemented, cg_timer shall expire synchronously with the rising edge of TX_CLK as well as every one-eighth of the TX_CLK cycle time (see tolerance required for TX_CLK in 46.3.1.1). In the absence of XGMII, cg_timer shall expire every $3.2 \text{ ns} \pm 100 \text{ ppm}$.

Values: The condition cg_timer_done becomes true upon timer expiration.

Restart when: immediately after expiration; restarting the timer resets the condition cg_timer_done.

Duration: 3.2 ns nominal.

TX_CLK_timer

A continuous free-running timer. TX_CLK_timer shall expire synchronously with the rising edge of TX_CLK (see tolerance required for TX_CLK in 46.3.1.1).

Restart when: immediately after expiration.

Duration: 25.6 ns nominal.

The following timers are used only for the EEE capability:

rx_tq_timer

This timer is started when the PCS receiver enters the START_TQ_TIMER state. The timer terminal count is set to T_{QR} . When the timer reaches terminal count, it will set the rx_tq_timer_done = TRUE.

rx_tw_timer

This timer is started when the PCS receiver enters the RX_WAKE state. The timer terminal count shall not exceed the maximum value of T_{WR} in Table 127-7. When the timer reaches terminal count, it will set the rx_tw_timer_done = TRUE.

rx_wf_timer

This timer is started when the PCS receiver enters the RX_WTF state, indicating that the receiver has encountered a wake time fault. The rx_wf_timer allows the receiver an additional period in which to synchronize or return to the quiescent state before a link failure is indicated. The timer terminal count is set to T_{WTF} . When the timer reaches terminal count, it will set the rx_wf_timer_done = TRUE.

tx_tq_timer

This timer is started when the PCS transmitter enters the TX_QUIET state. The timer terminal count is set to T_{QL} . When the timer reaches terminal count, it will set the tx_tq_timer_done = TRUE.

tx_tr_timer

This timer is started when the PCS transmitter enters the TX_REFRESH state. The timer terminal count is set to T_{UL} . When the timer reaches terminal count, it will set the tx_tr_timer_done = TRUE.

tx_ts_timer

This timer is started when the PCS transmitter enters the TX_SLEEP state. The timer terminal count is set to T_{SL} . When the timer reaches terminal count, it will set the tx_ts_timer_done = TRUE.

127.2.7.2 State diagrams

127.2.7.2.1 Word Encode and Word-to-Octets

The Word Encode process (see 127.2.5.2) and Word-to-Octets process (see 127.2.5.3) are merged into one state diagram depicted in Figure 127–4, including compliance with the associated state variables as specified in 127.2.7.1.

The Word Encode process continuously maps the four XGMII lanes to four 2.5GPPII symbols via the WENCODE function in the TX_XGMII_LO and TX_XGMII_HI states. The four 2.5GPPII symbols are then serialized and output one at a time by the Word-to-Octets process. The presentation of the 2.5GPPII symbols are synchronized to the PCS transmit process such that the 2.5GPPII index 0 and 2 symbols are presented to the PCS transmit process which will result in the corresponding ordered set to be output to the PMA when the variable tx_even is TRUE and index 1 and 3 variables when tx_even is FALSE.

127.2.7.2.2 Transmit

The PCS Transmit process is depicted in two state diagrams: PCS Transmit ordered set and PCS Transmit code-group. The PCS shall implement its Transmit process as depicted in Figure 127–5 and Figure 127–6, including compliance with the associated state variables as specified in 127.2.7.1.

The Transmit ordered set process continuously sources ordered sets to the Transmit code-group process. Upon the assertion of tp_en by the 2.5GPPII, the SPD ordered set is sourced. Following the SPD, /D/ code-groups are sourced until tp_en is deasserted. Following the de-assertion of tp_en, EPD ordered sets are sourced. If tp_en and tp_er are both deasserted, the /R/ ordered set may be sourced, after which the sourcing of /I/ is resumed. If, while tp_en is asserted, the tp_er signal is asserted, the /V/ ordered set is sourced except when the SPD ordered set is selected for sourcing. If the 2.5GPPII indicates sequence, then /Q/ ordered sets are sourced. If the optional EEE is enabled and the 2.5GPPII indicates low power idles, then /LI/ ordered sets are sourced.

The Transmit code-group process continuously sources tx_code-group<9:0> to the PMA based on the ordered sets sourced to it by the Transmit ordered set process. The Transmit code-group process determines the proper code-group to source based on even/odd-numbered code-group alignment, running disparity requirements, and ordered set format.

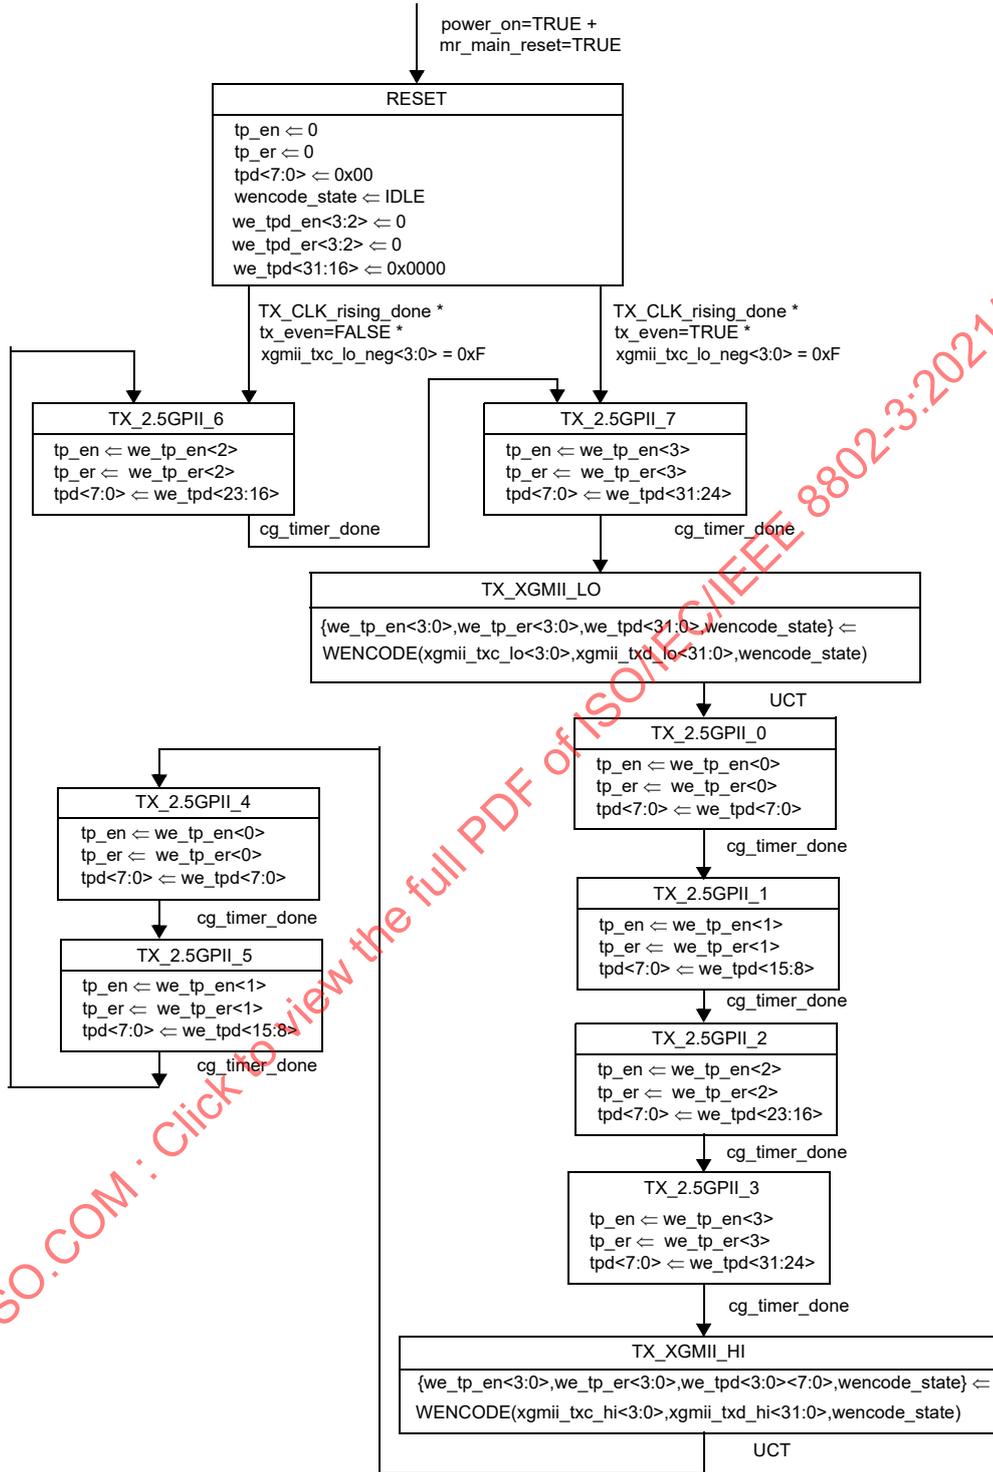


Figure 127-4—PHY TX control state diagram

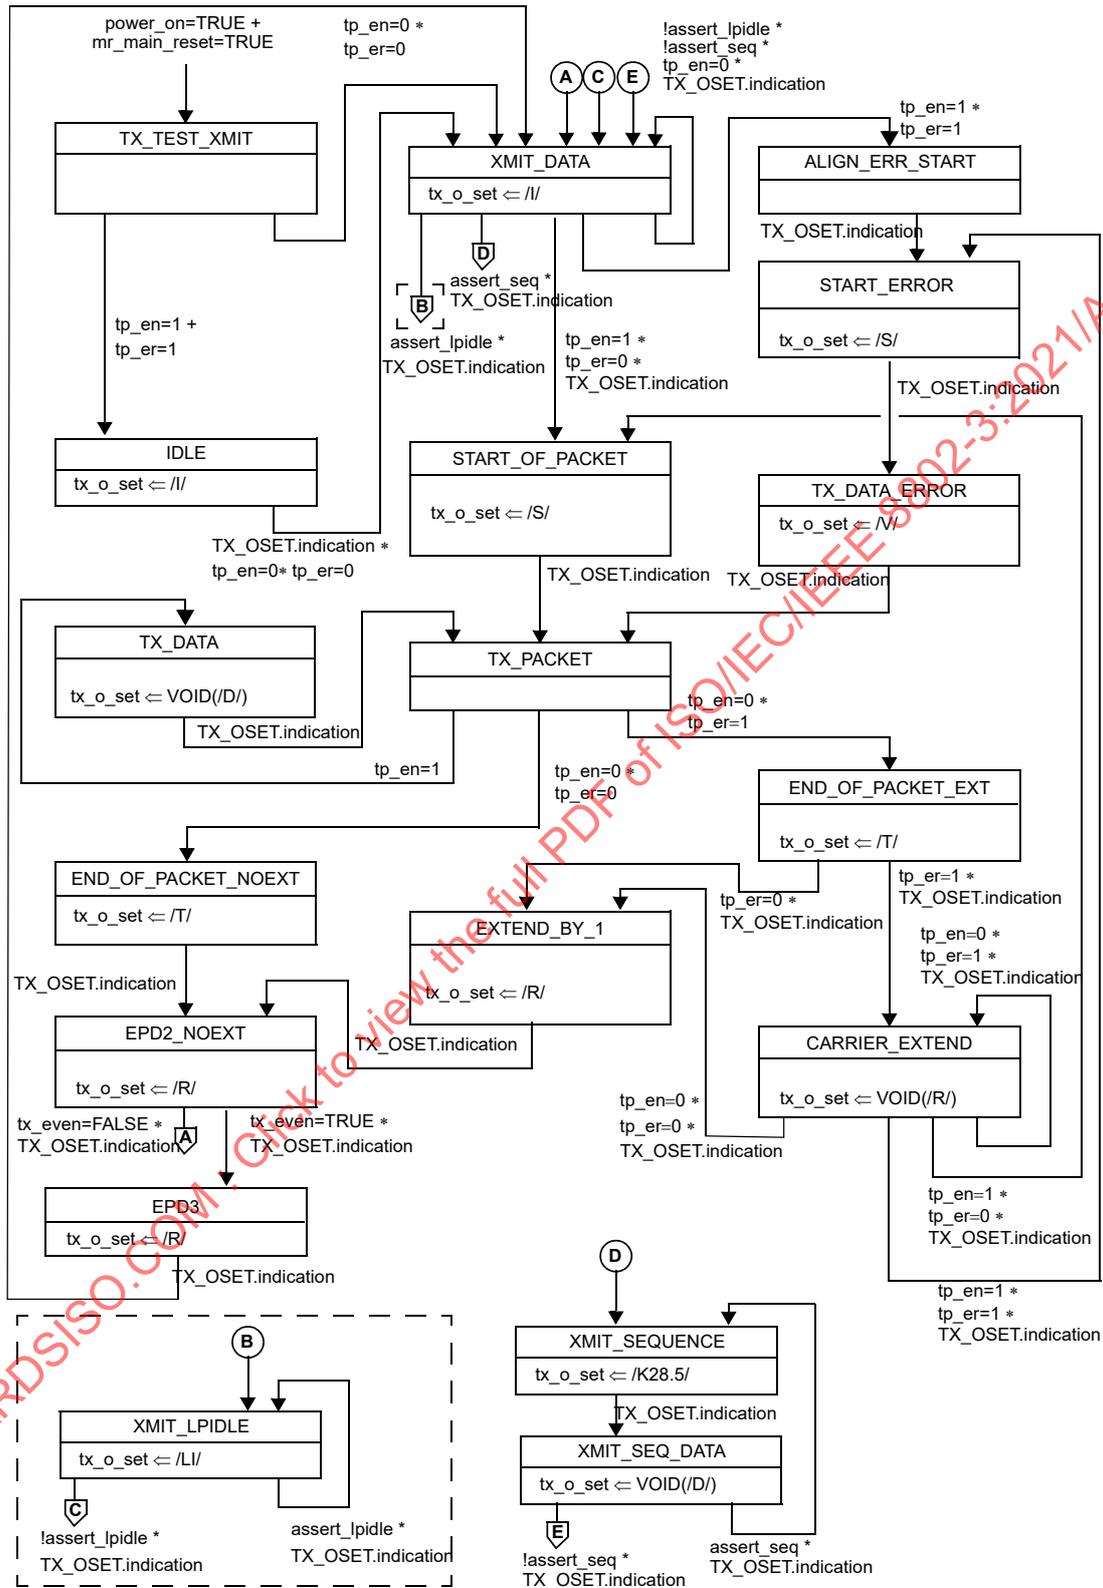


Figure 127-5—PCS transmit ordered set state diagram

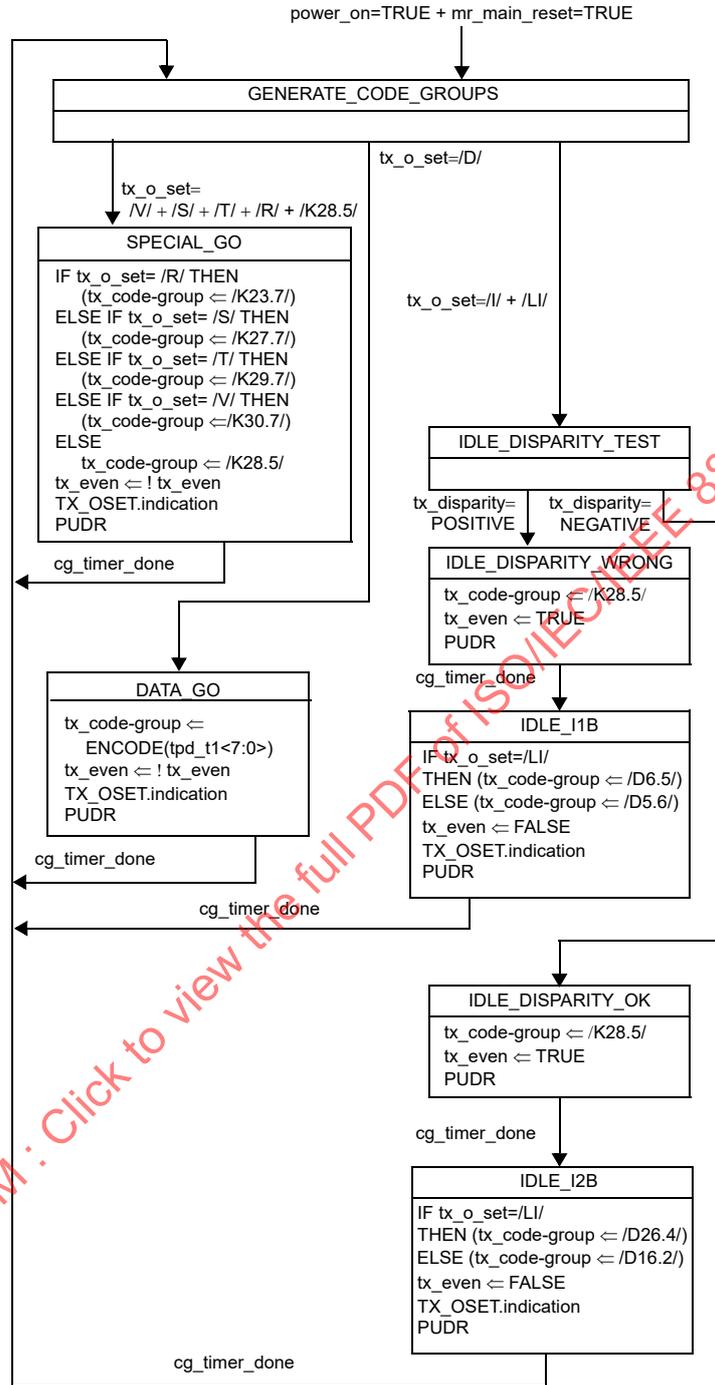


Figure 127-6—PCS transmit code-group state diagram

127.2.7.2.3 Synchronization

The PCS shall implement the Synchronization process as depicted in Figure 127–7 including compliance with the associated state variables as specified in 127.2.7.1. The Synchronization process is responsible for determining whether the underlying receive channel is ready for operation. Failure of the underlying channel typically causes the PMA client to suspend normal actions.

A receiver that is in the LOSS_OF_SYNC state and that has acquired bit synchronization attempts to acquire code-group synchronization via the Synchronization process. Code-group synchronization is acquired by the detection of three ordered sets containing commas in their leftmost bit positions without intervening invalid code-group errors. Upon acquisition of code-group synchronization, the receiver enters the SYNC_ACQUIRED_1 state. Acquisition of synchronization ensures the alignment of multi-code-group ordered sets to even-numbered code-group boundaries.

Once synchronization is acquired, the Synchronization process tests received code-groups in sets of four code-groups and employs multiple sub-states, to move between the SYNC_ACQUIRED_1 and LOSS_OF_SYNC states.

For EEE capability the relationship between sync_status and code_sync_status is given by the definition of the sync_status variable in 127.2.7.1.3; otherwise sync_status is identical to code_sync_status.

127.2.7.2.4 Receive

The PCS shall implement its Receive process as depicted in Figure 127–8a and Figure 127–8b, including compliance with the associated state variables as specified in 127.2.7.1. The PCS shall implement Figure 127–8c if the optional EEE is present and enabled.

The PCS Receive process continuously passes rpd<7:0> and sets the rp_dv and rp_er signals to the 2.5GPII based on the received code-group from the PMA.

127.2.7.2.5 Octets-to-Word and Decode

The Octets-to-Word process and Word decode process are merged into one state diagram depicted in Figure 127–9, including compliance with the associated state variables as specified in 127.2.7.1.

The Octets-to-Word process continuously queues the incoming 2.5GPII symbols from the PCS receive process and outputs four 2.5GPII symbols at a time using the WALIGN function. Symbols may be dropped or idles symbols added by the WALIGN function. The Word Decode process continuously maps the four 2.5GPII symbols and presents them on the XGMII using the WDECODE function in the RX_XGMII state.

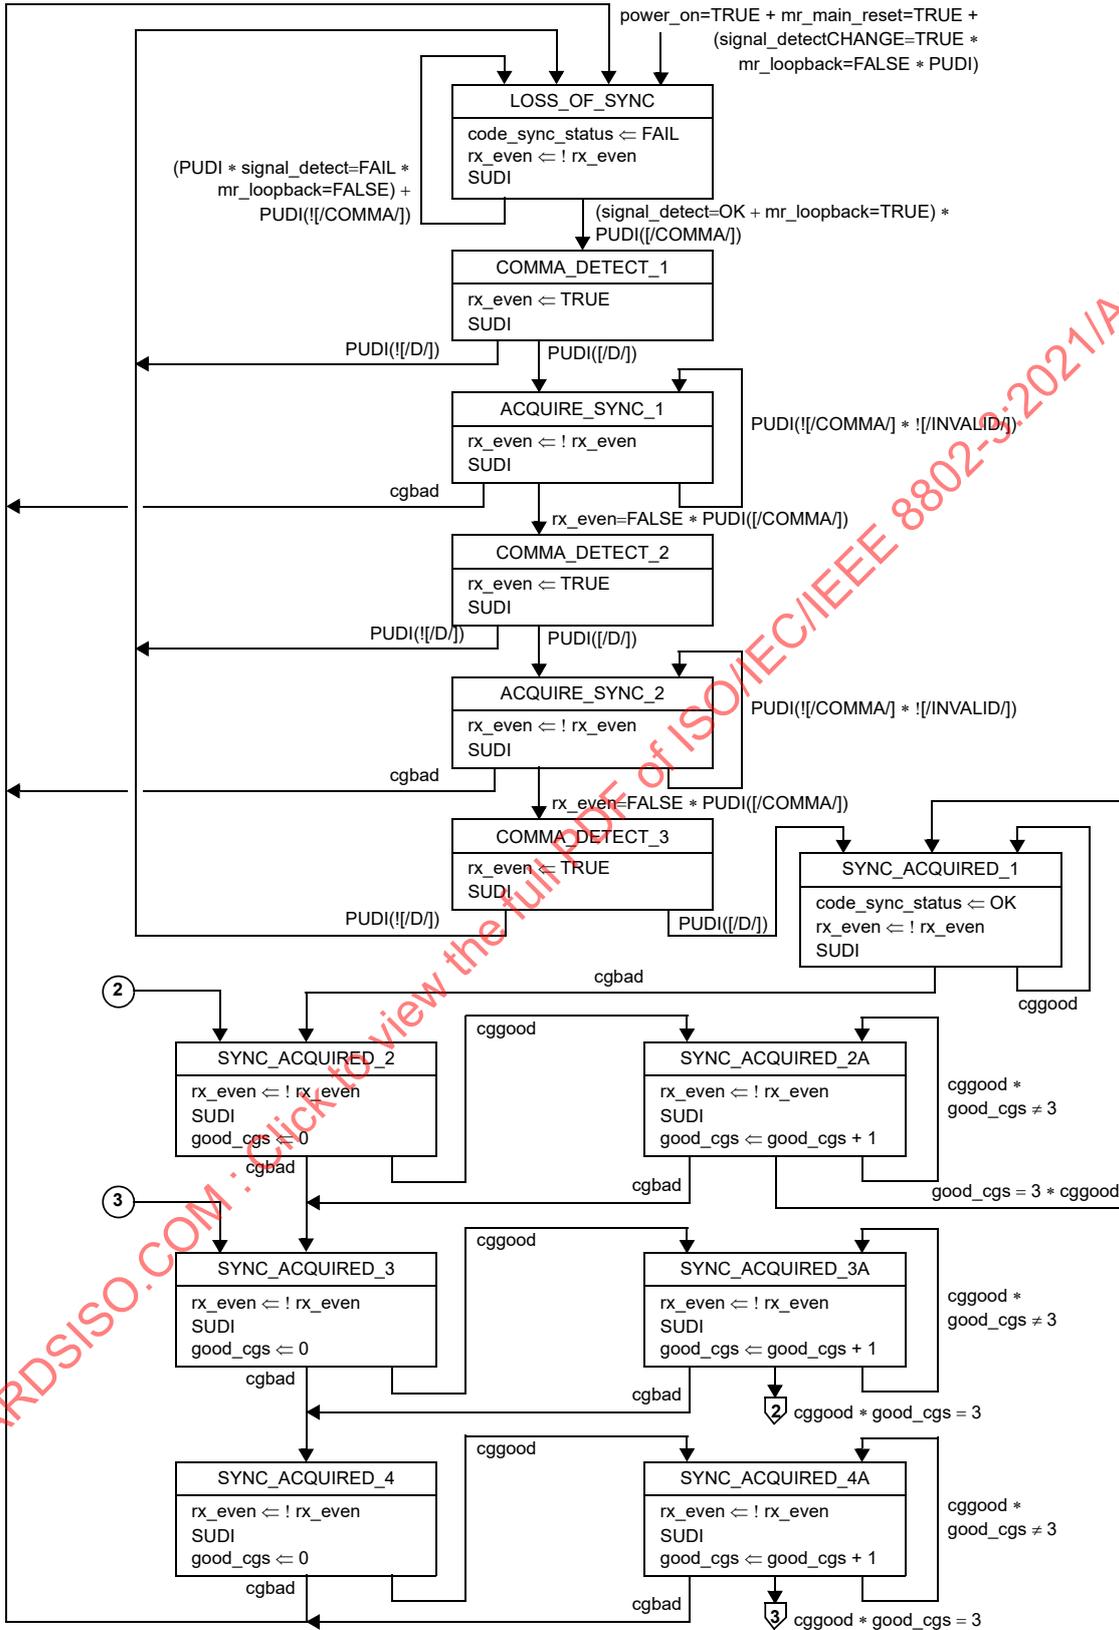
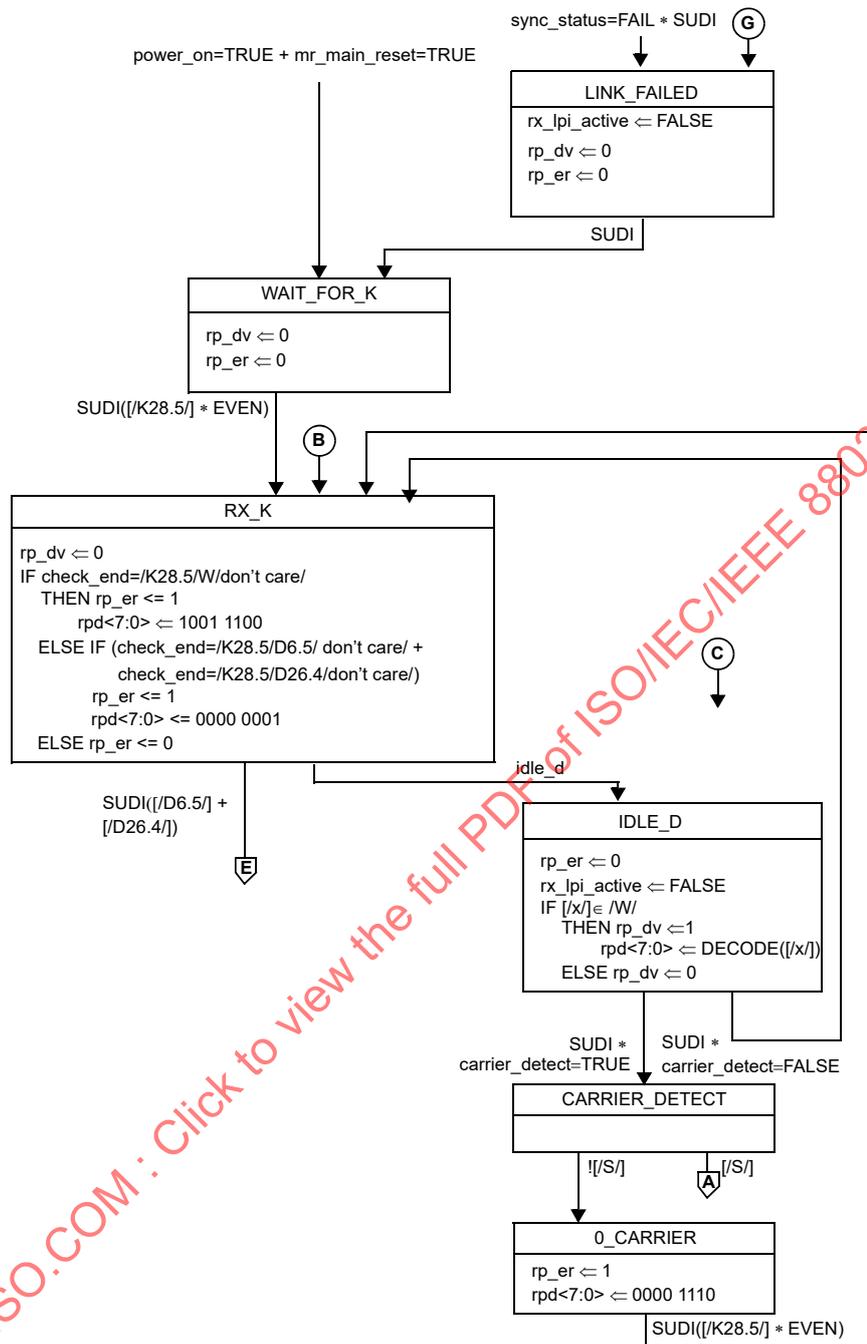


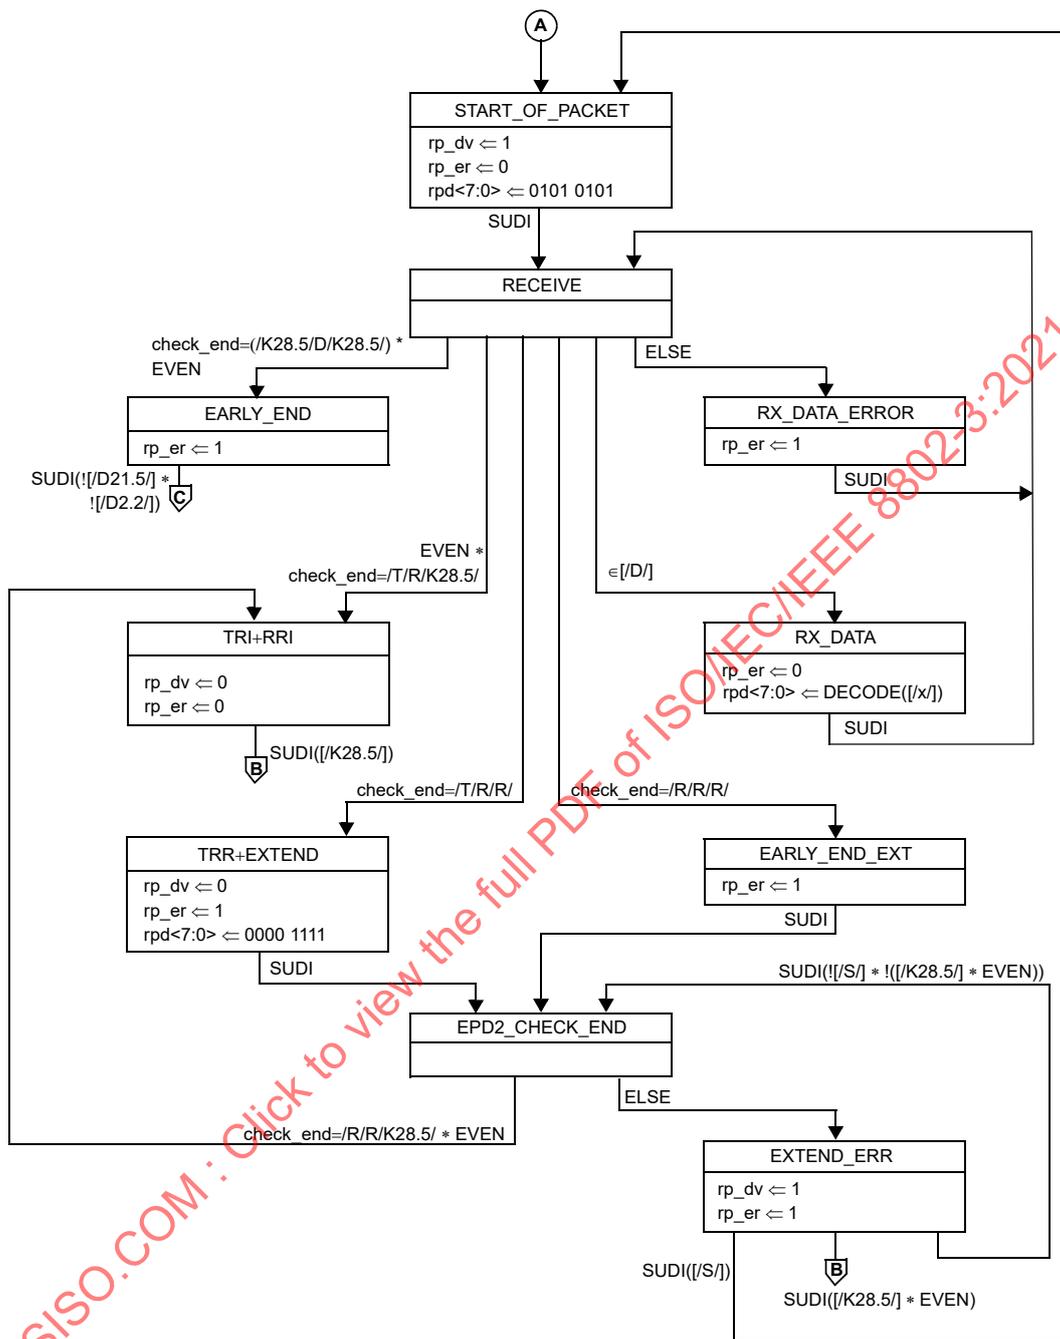
Figure 127-7—Synchronization state diagram



NOTE 1—Outgoing arcs leading to labeled polygons flow off page to corresponding incoming arcs leading from labeled circles on Figure 127–8b and Figure 127–8c, and vice versa.

NOTE 2—The transitions from the CARRIER_DETECT state is a test against the codegroup obtained from the SUDI that caused the transition to the CARRIER_DETECT state.

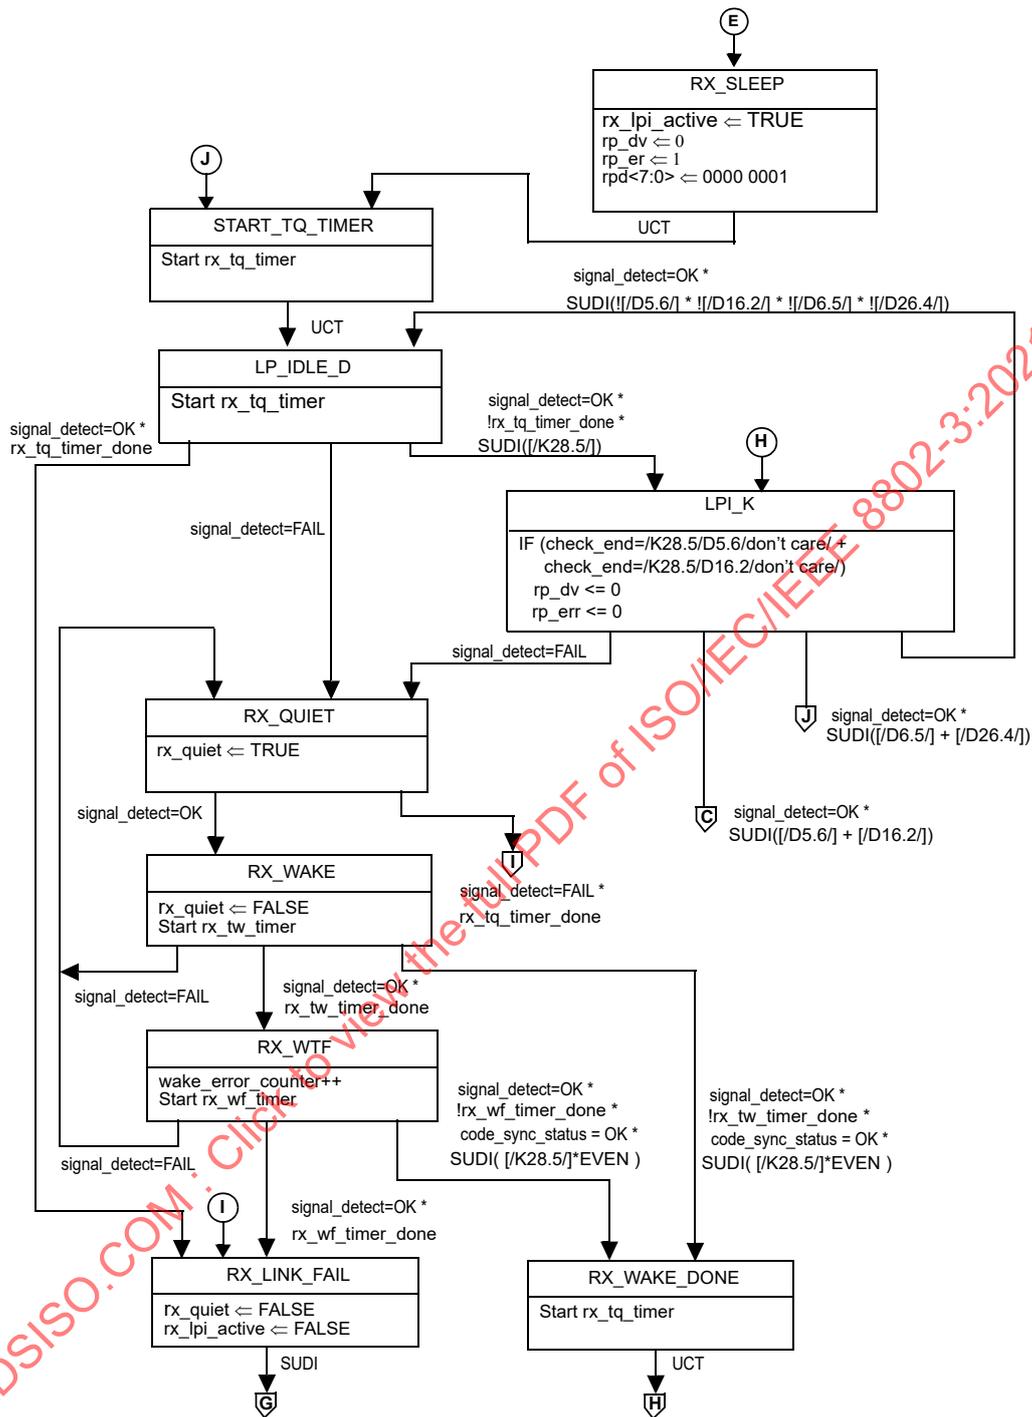
Figure 127–8a—PCS receive state diagram, part a



NOTE 1—Outgoing arcs leading to labeled polygons flow off page to corresponding incoming arcs leading from labeled circles on Figure 127–8a, and vice versa.

NOTE 2—In the transition from RECEIVE to RX_DATA state the transition condition is a test against the code-group obtained from the SUDI that caused the transition to RECEIVE state.

Figure 127–8b—PCS receive state diagram, part b



NOTE—Outgoing arcs leading to labeled polygons flow off page to corresponding incoming arcs leading from labeled circles on Figure 127–8a, and vice versa.

Figure 127–8c—PCS Receive state diagram, part c
 (only required for the optional EEE capability)

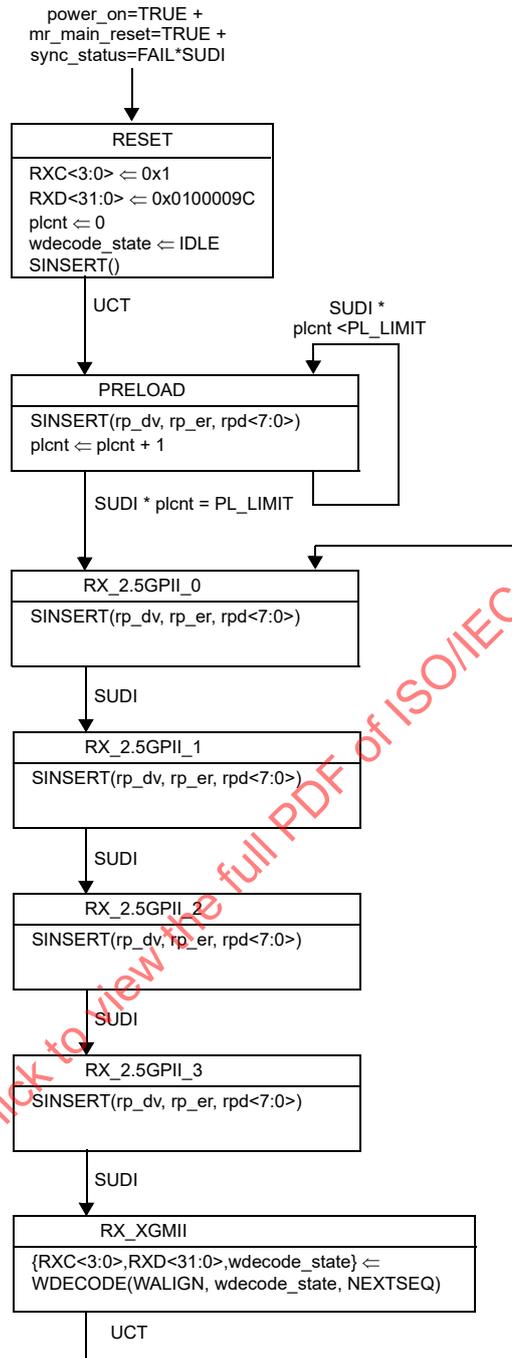


Figure 127-9—Octets-to-Word and Decode state diagram

127.2.7.2.6 LPI state diagram

A PCS that supports the EEE capability shall implement the LPI transmit process as shown in Figure 127–10. The transmit LPI state diagram controls tx_quiet, which disables the transmitter when true.

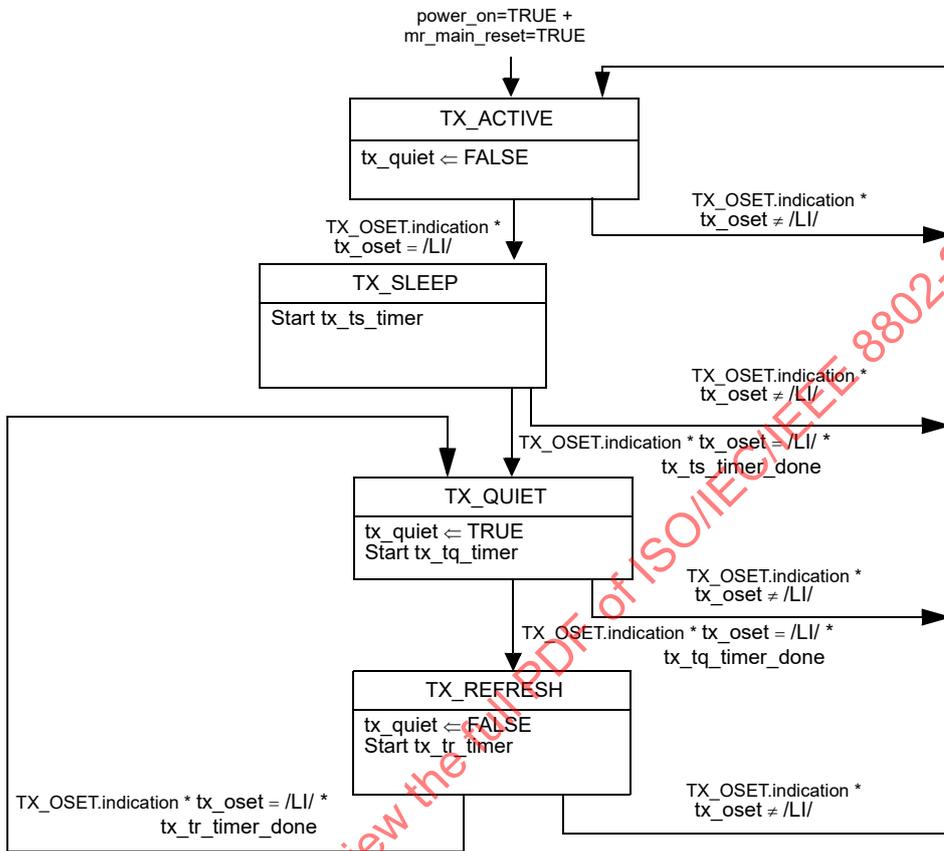


Figure 127–10—LPI Transmit state diagram

The timer values for these state diagrams are shown in Table 127–6 for transmit and Table 127–7 for receive.

Table 127–6—Transmitter LPI timing parameters

Parameter	Description	Min	Max	Units
T _{SL}	Local Sleep Time from entering the TX_SLEEP state to when tx_quiet is set to TRUE	19.9	20.1	μs
T _{QL}	Local Quiet Time from when tx_quiet is set to TRUE to entry into the TX_REFRESH state	2.5	2.6	ms
T _{UL}	Local Refresh Time from entry into the TX_REFRESH state to entry into the TX_QUIET state	19.9	20.1	μs

Table 127–7—Receiver LPI timing parameters

Parameter	Description	Min	Max	Units
T _{QR}	The time the receiver waits for signal detect to be set to OK while in the LP_IDLE_D, LPI_K and RX_QUIET states before asserting a rx_fault	3	4	ms
T _{WR}	Time the receiver waits in the RX_WAKE state before indicating a wake time fault (WTF)		11	μs
T _{WTF}	Wake time fault recovery time		1	ms

127.2.7.2.7 LPI status and management

For EEE capability, the PCS indicates to the management system that LPI is currently active in the receive and transmit directions using the status variables shown in Table 36-10.

127.3 Physical Medium Attachment (PMA) sublayer

127.3.1 Service Interface

The PMA provides a Service Interface to the PCS. These services are described in an abstract manner and do not imply any particular implementation. The PMA Service Interface supports the exchange of code-groups between PCS entities. The PMA converts code-groups into bits and passes these to the PMD, and vice versa. It also generates an additional status indication for use by its client.

The following primitives are defined:

```
PMA_UNITDATA.request(tx_code-group<9:0>)
PMA_UNITDATA.indication(rx_code-group<9:0>)
```

127.3.1.1 PMA_UNITDATA.request

This primitive defines the transfer of data (in the form of code-groups) from the PCS to the PMA. PMA_UNITDATA.request is generated by the PCS Transmit process.

127.3.1.1.1 Semantics of the service primitive

```
PMA_UNITDATA.request(tx_code-group<9:0>)
```

The data conveyed by PMA_UNITDATA.request is the tx_code-group<9:0> parameter defined in 127.2.7.1.3.

127.3.1.1.2 When generated

The PCS continuously sends, at a nominal rate of 312.5 MHz, tx_code-group<9:0> to the PMA.

127.3.1.1.3 Effect of receipt

Upon receipt of this primitive, the PMA generates a series of ten PMD_UNITDATA.request primitives, requesting transmission of the indicated tx_bit to the PMD.

127.3.1.2 PMA_UNITDATA.indication

This primitive defines the transfer of data (in the form of code-groups) from the PMA to the PCS. PMA_UNITDATA.indication is used by the PCS Synchronization process.

127.3.1.2.1 Semantics of the service primitive

PMA_UNITDATA.indication(rx_code-group<9:0>)

The data conveyed by PMA_UNITDATA.indication is the rx_code-group<9:0> parameter defined in 127.2.7.1.3.

127.3.1.2.2 When generated

The PMA continuously sends one rx_code-group<9:0> to the PCS corresponding to the receipt of each code-group aligned set of ten PMD_UNITDATA.indication primitives received from the PMD. The nominal rate of the PMA_UNITDATA.indication primitive is 312.5 MHz, as governed by the recovered bit clock.

127.3.1.2.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMA sublayer.

127.3.2 Functions within the PMA

Figure 127–2 and Figure 127–3 depict the mapping of the four octet-wide data path of the XGMII to the ten-bit code-groups of the PMA Service Interface and on to the serial PMD Service Interface. The PMA comprises the PMA Transmit and PMA Receive processes for 2.5GBASE-X.

The PMA Transmit process serializes tx_code-groups into tx_bits and passes them to the PMD for transmission on the underlying medium, according to Figure 127–3. Similarly, the PMA Receive process deserializes rx_bits received from the PMD according to Figure 127–3. The PMA continuously conveys ten bit code-groups to the PCS, independent of code-group alignment. After code-group alignment is achieved, based on comma detection, the PCS converts code-groups into XGMII data octets, according to 127.2.7.2.4 and 127.2.7.2.5.

The proper alignment of a comma used for code-group synchronization is depicted in Figure 127–3.

127.3.2.1 Data delay

The PMA maps a nonaligned one-bit data path from the PMD to an aligned, ten-bit-wide data path to the PCS, on the receive side. Logically, received bits must be buffered to facilitate proper code-group alignment. These functions necessitate an internal PMA delay of at least ten bit times. In practice, code-group alignment may necessitate even longer delays of the incoming rx_bit stream.

127.3.2.2 PMA transmit function

The PMA Transmit function passes data unaltered (except for serializing) from the PCS directly to the PMD. Upon receipt of a PMA_UNITDATA.request primitive, the PMA Transmit function shall serialize the ten bits of the tx_code-group<9:0> parameter and transmit them to the PMD in the form of ten successive PMD_UNITDATA.request primitives, with tx_code-group<0> transmitted first, and tx_code-group<9> transmitted last.

127.3.2.3 PMA receive function

The PMA Receive function passes data unaltered (except for deserializing and possible code-group slipping upon code-group alignment) from the PMD directly to the PCS. Upon receipt of ten successive PMD_UNITDATA.indication primitives, the PMA shall assemble the ten received rx_bits into a single ten-bit value and pass that value to the PCS as the rx_code-group<9:0> parameter of the primitive PMA_UNITDATA.indication, with the first received bit installed in rx_code-group<0> and the last received bit installed in rx_code-group<9>. An exception to this operation is specified in 127.3.2.4.

127.3.2.4 Code-group alignment

In the event the PMA sublayer detects a comma+ within the incoming rx_bit stream, it may realign its current code-group boundary, if necessary, to that of the received comma+ as shown in Figure 36-3. This process is referred to in this document as code-group alignment. During the code-group alignment process, the PMA sublayer may delete or modify up to four, but shall delete or modify no more than four, ten-bit code-groups in order to align the correct receive clock and code-group containing the comma+. This process is referred to as code-group slipping.

In addition, the PMA sublayer is permitted to realign the current code-group boundary upon receipt of a comma-pattern.

127.3.3 Loopback mode

Loopback mode shall be provided, as specified in this subclause, by the transmitter and receiver of a device as a test function to the device. When Loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. A device is explicitly placed in Loopback mode (i.e., Loopback mode is not the normal mode of operation of a device). The method of implementing Loopback mode is not defined by this standard.

NOTE—Loopback mode may be implemented either in the parallel or the serial circuitry of a device.

127.3.3.1 Receiver considerations

A receiver may be placed in Loopback mode. Entry into or exit from Loopback mode may result in a temporary loss of synchronization.

127.3.3.2 Transmitter considerations

A transmitter may be placed in Loopback mode. While in Loopback mode, the transmitter output is not defined.

127.3.4 Test functions

A limited set of test functions may be provided as an implementation option for testing of the transmitter function or for testing of an attached receiver.

Some test functions that are not defined by this standard may be provided by certain implementations. Compliance with the standard is not affected by the provision or exclusion of such functions by an implementation. The patterns described in Annex 36A may be used for 2.5GBASE-X except the nominal bit rate is 2.5 times faster and any references to the GMII applies to the XGMII.

A typical test function is the ability to transmit invalid code-groups within an otherwise valid PHY bit stream. Certain invalid PHY bit streams may cause a receiver to lose word and/or bit synchronization. See

ANSI X3.230-1994 [B21] (FC-PH), subclause 5.4, for a more detailed discussion of receiver and transmitter behavior under various test conditions.

127.3.4.1 PMA PRBS9 test pattern (optional)

The PMA may optionally generate a PRBS9 test pattern in the transmit direction.

The ability to generate the test pattern is indicated by the PRBS9_Tx_generator_ability status variable, which, if a Clause 45 MDIO is implemented, is accessible through bit 1.1500.5 (see 45.2.1.140).

If supported, when send Tx PRBS9 test-pattern mode is enabled by the PRBS9_enable and PRBS_Tx_gen_enable control variables, the PMA shall generate a PRBS9 pattern (as defined in footnote a of Table 68–6) toward the service interface below the PMA via the PMA_UNITDATA.request primitive. If a Clause 45 MDIO is implemented, the PRBS9_enable and PRBS_Tx_gen_enable control variables are accessible through bits 1.1501.6 and 1.1501.3 (see 45.2.1.141). When send Tx PRBS9 test-pattern mode is disabled, the PMA returns to normal operation.

Note that PRBS9 is intended to be checked by external test gear, and no PRBS9 checking function is provided within the PMA.

127.4 Compatibility considerations

There is no requirement for a compliant device to implement or expose any of the interfaces specified for the PCS or PMA. Implementations of a XGMII shall comply with the requirements as specified in Clause 46.

127.5 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) also 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 backplane channel and concatenation of devices.

The sum of transmit and receive delay contributed by the 2.5GBASE-X PCS and PMA shall be no more than 768 bit-times.

127.6 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 71.9.

**127.7 Protocol implementation conformance statement (PICS) proforma for
 Clause 127, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA)
 sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X⁵**

127.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 127, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

127.7.2 Identification

127.7.2.1 Implementation identification

Supplier (Note 1)	
Contact point for inquiries about the PICS (Note 1)	
Implementation Name(s) and Version(s) (Notes 1 and 3)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) (Note 2)	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

127.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3cb-2018, Clause 127, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 2.5 Gb/s 8B/10B 2.5GBASE-X
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3cb-2018.)	
Date of Statement	

⁵Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

127.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
*XGE	XGMII compatibility interface	127.1.4, Clause 46	Compatibility interface is supported	O	Yes [] No []
MD	MDIO	Clause 45, 49.2.14	Registers and interface supported	O	Yes [] No []
*LPI	Implementation of LPI	Clause 78		O	Yes [] No []
LNKS	Implementation of PCS Link Status Signalling	127.2.6.6		O	Yes [] No []

127.7.4 PICS proforma tables for the PCS and PMA sublayer, type 2.5GBASE-X

127.7.4.1 PCS

Item	Feature	Subclause	Value/Comment	Status	Support
PCS1	Word-to-Octet Synchronization	127.2.5.3		M	Yes []
PCS2	A transition of a 2.5GPll idle symbol to data or error symbol	127.2.5.4	Place data or error symbol on index 0	M	Yes []
PCS3	A transition of a 2.5GPll idle symbol to LPI symbol	127.2.5.4	Place LPI symbol on either index 0 or 2	M	Yes []
PCS4	A transition of a 2.5GPll LPI symbol to idle symbol	127.2.5.4	Place idle symbol on either index 0 or 2	M	Yes []
PCS5	Q or Fsig set start	127.2.5.4	Always occur on index 0	M	Yes []
PCS6	Deficit Idle Count limits	127.2.5.4	Bounded to a min of 0, max of 3	M	Yes []
PCS7	2.5GBASE-X 8B/10B requirement	127.2.6		M	Yes []
PCS8	PCS Transmit Process	127.2.7.2.2	State diagrams Figure 127–5, and Figure 127–6 including compliance with variables in 127.2.7.1	M	Yes []
PCS9	PCS Synchronization	127.2.7.2.3	State diagram Figure 127–7, including compliance with variables in 127.2.7.1	M	Yes []
PCS10	PCS Receive Process	127.2.7.2.4	State diagrams Figure 127–8a, and Figure 127–8b including compliance with variables in 127.2.7.1	M	Yes []

127.7.4.2 Code-group functions

Item	Feature	Subclause	Value/Comment	Status	Support
CG1	Word-to-Octets process shall be synchronized to the PCS transmit process	127.2.5.3	Ordered set to be output to the PMA	M	Yes []
CG2	A transition of a 2.5GPPII idle symbol to a data or error symbol	127.2.5.4	Data or error symbol on index 0	M	Yes []
CG3	A transition of a 2.5GPPII idle symbol to a LPI symbol	127.2.5.4	LPI symbol on either index 0 or 2	M	Yes []
CG4	A transition of a 2.5GPPII LPI symbol to an idle symbol	127.2.5.4	Idle symbol on either index 0 or 2	M	Yes []
CG5	The start of Q or Fsig set	127.2.5.4	Always occur on index 0	M	Yes []
CG6	Deficit Idle Count (DIC) limits	127.2.5.4	Bounded to a min of 0, max of 3	M	Yes []
CG7	2.5GBASE-X 8B/10B requirement	127.2.6		M	Yes []

127.7.4.3 EEE

Item	Feature	Subclause	Value/Comment	Status	Support
EEE1	Word-to-Octet Transitions - idle to data or error	127.2.5.4	Place LPI symbol on index 0 or 2	LPI:M	Yes [] N/A []
EEE2	rx_tw_timer terminal count	127.2.7.1.7	Timer terminal count not to exceed TWR in Table 127-7	LPI:M	Yes [] N/A []
EEE3	PCS Receive Process when EEE is present	127.2.7.2.4	State diagram Figure 127-8c including compliance with variables in 127.2.7.1	LPI:M	Yes [] N/A []
EEE4	PCS that supports EEE capability	127.2.7.2.6	Implement LPI transmit process shown in Figure 127-10	LPI:M	Yes [] N/A []

127.7.4.4 PMA functions

Item	Feature	Subclause	Value/Comment	Status	Support
PMA1	Transmit function	127.3.2.2	Serialize 10 bits of tx_code-group	M	Yes []
PMA2	Receive function	127.3.2.3	PMA assembles received rx_bits into a single ten-bit value and passes it to PCS	M	Yes []
PMA3	Code-group alignment	127.3.2.4	PMA sublayer shall delete or modify ≤ 4 , ten-bit code groups	M	Yes []
PMA4	Loopback mode	127.3.3		M	Yes []

127.7.4.5 Compatibility considerations

Item	Feature	Subclause	Value/Comment	Status	Support
CC1	XGMII	127.4, 46		XGE:M	Yes [] N/A []
CC2	Transmit and Receive PCS/PMA delay constraints	127.5	Delay contributions by 2.5GBASE-X PCS and PMA sublayers ≤ 768 bit times	M	Yes []
CC3	Environmental specifications	127.6	Conform to 71.9	M	Yes []

128. Physical Medium Dependent sublayer and baseband medium, type 2.5GBASE-KX

128.1 Overview

This clause specifies the 2.5GBASE-KX PMD and baseband medium. When forming a complete PHY, a PMD shall be combined with the appropriate sublayers (see Table 128–1), and with the management functions that are optionally accessible through the management interface defined in Clause 45. References to the MDI (Media Dependent Interface) should be considered to be TP1 for the transmitter and TP4 for the receiver, as measurement points.

Table 128–1—Physical Layer clauses associated with the 2.5GBASE-KX PMD

Associated clause	2.5GBASE-KX
46—XGMII ^a	Optional
73—Auto-Negotiation for Backplane	Required
78—EEE	Optional
127—2.5GBASE-X PCS/PMA	Required

^aThe XGMII is an optional interface. However, if the XGMII is not implemented, a conforming implementation must behave functionally as though the RS and XGMII were present.

A 2.5GBASE-KX PHY with the optional Energy Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

128.2 Physical Medium Dependent (PMD) service interface

The following specifies the services provided by the 2.5GBASE-KX PMDs. These PMD sublayer service interfaces are described in an abstract manner and do not imply any particular implementation.

The PMD Service Interface supports the exchange of bit streams representing 2.5GBASE-X 8B/10B encoded data between the PMA and PMD entities. The PMD translates the serialized data of the PMA to and from signals suitable for the specified medium.

The following primitives are defined:

```

PMD_UNITDATA.request(tx_bit)
PMD_UNITDATA.indication(rx_bit)
PMD_SIGNAL.indication(SIGNAL_DETECT)

```

If EEE is supported, the following primitives are also defined on the PMD Service Interface:

```

PMD_RXQUIET.request(rx_quiet)
PMD_TXQUIET.request(tx_quiet)

```

These messages affect the PCS variables as described in 127.2.7.1.6.

128.2.1 PMD_UNITDATA.request

This primitive defines the transfer of a serial data stream from the PMA to the PMD.

128.2.1.1 Semantics of the service primitive

PMD_UNITDATA.request(tx_bit)

The data conveyed by PMD_UNITDATA.request is a continuous stream of bits. The tx_bit parameter can take one of two values: ONE or ZERO.

128.2.1.2 When generated

The PMA continuously sends the appropriate stream of bits to the PMD for transmission on the medium, at the signaling speed for 2.5GBASE-KX PMD types specified in Table 128–4.

128.2.1.3 Effect of receipt

Upon receipt of this primitive, the PMD converts the specified stream of bits into the appropriate signals on the MDI.

128.2.2 PMD_UNITDATA.indication

This primitive defines the transfer of a serial data stream from the PMD to the PMA.

128.2.2.1 Semantics of the service primitive

PMD_UNITDATA.indication(rx_bit)

The data conveyed by PMD_UNITDATA.indication is a continuous stream of bits. The rx_bit parameter can take one of two values: ONE or ZERO.

128.2.2.2 When generated

The PMD continuously sends a stream of bits to the PMA corresponding to the signals received from the MDI.

128.2.2.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

128.2.3 PMD_SIGNAL.indication

This primitive is generated by the PMD to indicate the status of the signal being received from the MDI.

128.2.3.1 Semantics of the service primitive

PMD_SIGNAL.indication(SIGNAL_DETECT)

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL, indicating whether the PMD is detecting signal at the receiver (OK) or not (FAIL). When SIGNAL_DETECT = FAIL, then rx_bit is undefined, but consequent actions based on PMD_UNITDATA.indication, where necessary, interpret rx_bit as a logic ZERO.

NOTE—SIGNAL_DETECT = OK does not guarantee that PMD_UNITDATA.indication(rx_bit) is known good. It is possible for a poor quality link to provide sufficient signal for a SIGNAL_DETECT = OK indication and still not meet the BER objective.

128.2.3.2 When generated

The PMD generates this primitive to indicate a change in the value of SIGNAL_DETECT. If the MDIO interface is implemented, then PMD_global_signal_detect shall be continuously set to the value of SIGNAL_DETECT.

128.2.3.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

128.2.4 PMD_RXQUIET.request

This primitive is generated by the PCS Receive Process when EEE is supported to indicate that the input signal is quiet and the PMA and PMD receiver may go into low power mode. When EEE is not supported, the primitive is never invoked, and the PMD behaves as if rx_quiet = FALSE.

128.2.4.1 Semantics of the service primitive

PMD_RXQUIET.request(rx_quiet)

The rx_quiet parameter takes on one of two values: TRUE or FALSE.

128.2.4.2 When generated

The PCS generates this primitive to request the appropriate PMD receive LPI state.

128.2.4.3 Effect of receipt

This variable is from the receive process of the PCS to control the power-saving function of the local PMD receiver. The 2.5GBASE-KX PHY receiver should put unused functional blocks into a low power state to save energy.

128.2.5 PMD_TXQUIET.request

This primitive is generated by the PCS Transmit Process when EEE is supported to indicate that the PMA and PMD transmit functions may go into a low power mode and to disable the PMD transmitter. See 128.6.5. When EEE is not supported, the primitive is never invoked, and the PMD behaves as if tx_quiet = FALSE.

128.2.5.1 Semantics of the service primitive

PMD_TXQUIET.request(tx_quiet)

The tx_quiet parameter takes on one of two values: TRUE or FALSE.

128.2.5.2 When generated

The PCS generates this primitive to request the appropriate PMD transmit LPI state.

128.2.5.3 Effect of receipt

This primitive affects operation of the PMD Transmit disable function as described in 128.6.5. The 2.5GBASE-KX PHY transmitter should put unused functional blocks into a lower power state to save energy.

128.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD shall support the AN service interface primitive AN_LINK.indication as defined in 73.9.

128.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must consider the delay maxima, and that network planners and administrators consider the delay constraints regarding the physical topology and concatenation of devices.

The sum of transmit and receive delays contributed by the 2.5GBASE-KX PMD and medium shall be no more than 256 bit times. It is assumed that the round-trip delay through the medium is less than or equal to 80 bit times.

128.5 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. If the MDIO is implemented, it shall map MDIO control variables to PMD control variables as shown in Table 128–2 and MDIO status variables to PMD status variables as shown in Table 128–3.

Table 128–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	Control register 1	1.0.15	PMD_reset
PMD Transmit Disable	1000BASE-KX/2.5GBASE-KX control register	1.160.0	PMD_transmit_disable

Table 128–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	Status register 1	1.1.7	PMD_fault
Transmit fault ability	1000BASE-KX/2.5GBASE-KX status register	1.161.13	PMD_Transmit_fault_ability
Receive fault ability	1000BASE-KX/2.5GBASE-KX status register	1.161.12	PMD_Receive_fault_ability
Transmit fault	1000BASE-KX/2.5GBASE-KX status register	1.161.11	PMD_transmit_fault
Receive fault	1000BASE-KX/2.5GBASE-KX status register	1.161.10	PMD_receive_fault
PMD transmit disable ability	1000BASE-KX/2.5GBASE-KX status register	1.161.8	PMD_transmit_disable_ability
Signal detect from PMD	1000BASE-KX/2.5GBASE-KX status register	1.161.0	PMD_signal_detect

128.6 PMD functional specifications

128.6.1 Link block diagram

For purposes of system conformance, the PMD sublayer is standardized at test points TP1 and TP4 as shown in Figure 128–1. The transmitter and receiver blocks include all off-chip components associated with the respective block. For example, external AC-coupling capacitors, if required, are to be included in the receiver block.

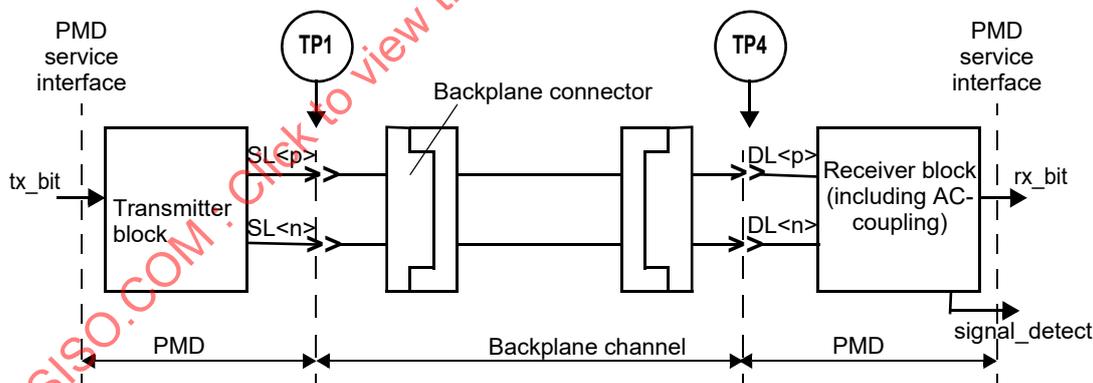


Figure 128–1—Link block diagram

The electrical path from the transmitter block to TP1, and from TP4 to the receiver block, will affect link performance and the measured values of electrical parameters used to verify conformance to this standard. Therefore, it is recommended that this path be carefully designed.

128.6.2 PMD transmit function

The PMD Transmit function shall convey the bits requested by the PMD service interface message `PMD_UNITDATA.request(tx_bit)` to the MDI according to the electrical specifications in 128.7.1. A positive output voltage of `SL<p>` minus `SL<n>` (differential voltage) shall correspond to `tx_bit = ONE`.

128.6.3 PMD receive function

The PMD Receive function shall convey the bits received at the MDI in accordance with the electrical specifications of 128.7.2 to the PMD service interface using the message `PMD_UNITDATA.indication(rx_bit)`. A positive input voltage of `DL<p>` minus `DL<n>` (differential voltage) shall correspond to `rx_bit = ONE`.

128.6.4 PMD signal detect function

For 2.5GBASE-KX operation PMD signal detect is mandatory if EEE is supported. When EEE is not supported, the PMD signal detect is optional for 2.5GBASE-KX, and its definition is beyond the scope of this specification. When PMD signal detect is not implemented, the value of `SIGNAL_DETECT` shall be set to OK for purposes of management and signaling of the primitive.

If EEE is supported, a local PMD signal detect function shall report to the PMD service interface using the message `PMD_SIGNAL.indication(SIGNAL_DETECT)`. This message is signaled continuously. For EEE, the `SIGNAL_DETECT` parameter can take on one of two values, OK or FAIL, indicating whether the PMD is detecting electrical energy at the receiver (OK) or not (FAIL). When `SIGNAL_DETECT = FAIL`, `PMD_UNITDATA.indication` is undefined. The signal energy from a compliant transmitter shall set `SIGNAL_DETECT` to OK within 750 ns when transitioning from LPI quiet to active and set `SIGNAL_DETECT` to FAIL within 750 ns when transitioning from active to LPI quiet.

128.6.5 PMD transmit disable function

The `PMD_transmit_disable` function is mandatory if EEE is supported and is otherwise optional. When this function is supported, it shall meet the following requirements.

- a) When the `PMD_transmit_disable` variable is set to ONE, this function shall turn off the transmitter such that it drives a constant level (i.e., no transitions) and does not exceed the maximum peak-to-peak differential output voltage specified in Table 128-4.
- b) If a `PMD_fault` (128.6.7) is detected, then the PMD may turn off the electrical transmitter.
- c) Loopback, as defined in 128.6.6, shall not be affected by `PMD_transmit_disable`.
- d) For EEE capability, the `PMD_transmit_disable` function shall turn off the transmitter after `tx_quiet` is asserted within the time and voltage level specified in 128.7.1.4. The `PMD_transmit_disable` function shall turn on the transmitter after `tx_quiet` is de-asserted within a time and voltage level specified in 128.7.1.4.

128.6.6 Loopback mode

Loopback mode shall be provided for the 2.5GBASE-KX PMA/PMD by the transmitter and receiver of a device as a test function to the device. When loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. Transmitter operation shall be independent of loopback mode. The method of implementing loopback mode within the PMD is not defined by this standard.

Control of the loopback function is specified in 45.2.1.1.5.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

128.6.7 PMD fault function

If the MDIO is implemented and the PMD has detected a local fault, the PMD shall set PMD_fault to ONE; otherwise, the PMD shall set PMD_fault to ZERO.

128.6.8 PMD transmit fault function

If the MDIO is implemented and the PMD has detected a local fault on the transmitter, the PMD shall set the PMD_transmit_fault variable to ONE; otherwise, the PMD shall set PMD_transmit_fault to ZERO.

128.6.9 PMD receive fault function

If the MDIO is implemented and the PMD has detected a local fault on the receiver, the PMD shall set the PMD_receive_fault variable to ONE; otherwise, the PMD shall set PMD_receive_fault to ZERO.

128.6.10 PMD LPI function

The PMD LPI function responds to PMD_TXQUIET and PMD_RXQUIET requests generated by the LPI transmit state diagram (See 127.2.7.2.7) and the LPI receive state diagram (see 49.2.13.3.1 and Figure 49-13). Implementation of the function is optional. EEE capabilities and parameters are advertised during the Backplane Auto-Negotiation as described in 45.2.7.15. The transmitter on the local device informs the link partner's receiver when to sleep, refresh, and wake. The local receiver's transitions are controlled by the link partner's transmitter and change independently from the local transmitter's states and transitions.

If EEE is supported, the PMD transmit function enters into a low power mode when tx_quiet is set to TRUE and exits when tx_quiet is set to FALSE. While tx_quiet is TRUE the PMD transmitter functional blocks should be deactivated to conserve energy. The PMD receive function enters into a low power mode when rx_quiet is set to TRUE and exits when rx_quiet is set to FALSE. While rx_quiet is TRUE the PMD receiver functional blocks should be deactivated to conserve energy.

The "Assert LPI" request at the XGMII is encoded in the transmitted symbols. Detection of LPI signaling in the received symbols is indicated as "Assert LPI" at the XGMII. Upon the detection of "Assert LPI" at the XGMII, an energy-efficient 2.5GBASE-KX PHY continues transmitting for a predefined period, then ceases transmission and deactivates transmit functions to conserve energy. The PHY periodically transmits during this quiet period to allow the remote PHY to refresh its receiver state (e.g., timing recovery, adaptive filter coefficients) and thereby track long-term variations in the timing of the link or the underlying channel characteristics. If, during the quiet or refresh periods, normal interframes resume at the XGMII, the PHY reactivates transmit functions and initiates transmission. This transmission will be detected by the remote PHY, causing it to also exit the LPI mode.

128.7 2.5GBASE-KX electrical characteristics

128.7.1 Transmitter characteristics

Transmitter characteristics at TP1 (see Figure 129–1) are summarized in Table 128–4 and detailed in 128.7.1.1 through 128.7.2.

Table 128–4—Transmitter characteristics for 2.5GBASE-KX

Parameter	Subclause reference	Value	Units
Signaling speed	128.7.1.3	3.125 ± 100 ppm	GBd
Peak-to-peak differential output voltage (max)	128.7.1.4	1200	mV
Peak-to-peak differential output voltage (min)	128.7.1.4	800	mV
Peak-to-peak differential output voltage (max) with TX disabled	128.6.5	30	mV
Common-mode voltage limits ^c	128.7.1.4	0 to +1.9	V
Differential output return loss (min)	128.7.1.5	See Equation (128–3)	dB
Common-mode output return loss (min)	128.7.1.6	See Equation (128–4)	dB
Transition time (20%–80%)	128.7.1.7	30 to 100	ps
Maximum Output jitter (peak-to-peak)			
Random jitter	128.7.1.9	0.2	UI
Deterministic jitter	128.7.1.9	0.12	UI
Duty Cycle Distortion ^a	128.7.1.9	0.035	UI
Total jitter ^b	128.7.1.9	0.32	UI

^aDuty Cycle Distortion is considered part of deterministic jitter distribution.
^bJitter is specified at BER 10⁻¹².
^cDefined with respect to signal ground as measured at Vcom in Figure 128–2.

128.7.1.1 Test fixtures

The test fixture of Figure 128–2, or its functional equivalent, is required for measuring the transmitter specifications described in 128.7.1, with the exception of return loss.

128.7.1.2 Test fixture characteristics

The differential load impedance applied to the transmitter output by the test fixture depicted in Figure 128–2 shall be 100 Ω. The differential return loss, in dB, of the test fixture shall meet the requirements of Equation (128–1) and Equation (128–2).

$$Return_loss(f) \geq 20 \tag{128-1}$$

for 100 MHz ≤ f < 1562.5 MHz and

$$Return_loss(f) \geq Return_loss_{min} = 20 - 20 \log_{10} \left(\frac{f}{1562.5 \text{ MHz}} \right) \tag{128-2}$$

for 1562.5 MHz ≤ f ≤ 2000 MHz.

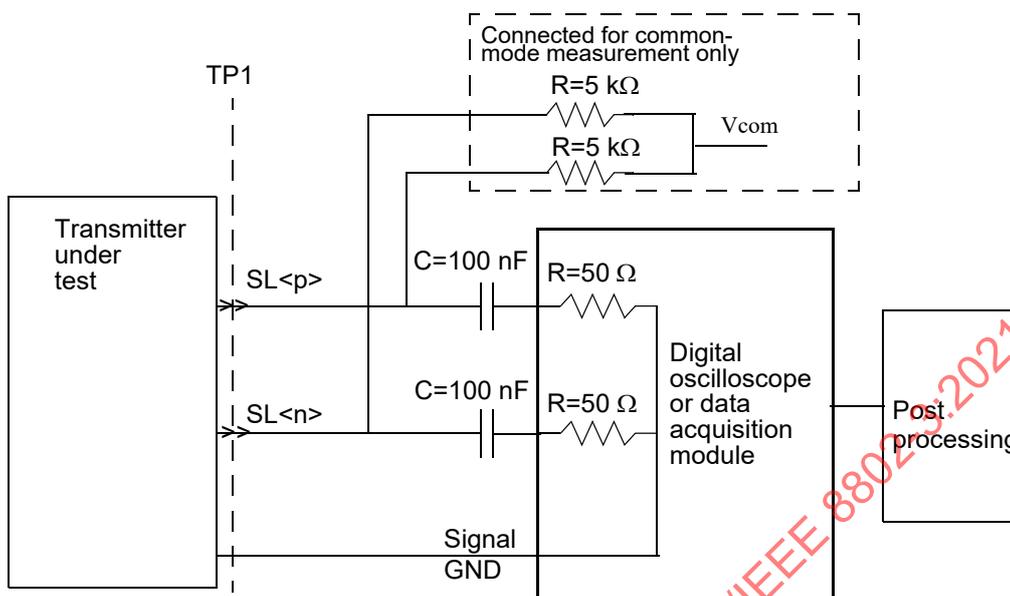


Figure 128-2—Transmit test fixture for 2.5GBASE-KX

128.7.1.3 Signaling speed

The 2.5GBASE-KX signaling speed shall be as specified in Table 128-4.

128.7.1.4 Output amplitude

The peak-to-peak differential output voltage shall be as specified in Table 128-4. See Figure 128-3. The differential output voltage test pattern is the test pattern specified in 48A.2. .

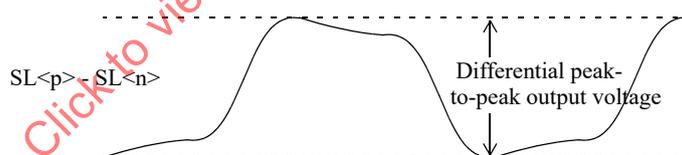


Figure 128-3—Transmitter peak-to-peak differential output voltage definition

NOTE—SL<p> and SL<n> are the positive and negative sides of the differential signal pair.

DC-referenced voltage levels are not defined since the receiver is AC-coupled. The common-mode voltage of SL<p> and SL<n>, with respect to signal ground, as measured at Vcom in Figure 128-2 is specified in Table 128-4.

For EEE capability, the transmitter shall meet the Table 128-4 requirement for Differential peak to peak output voltage when TX is disabled within 500 ns of tx_mode being set to QUIET and remain so while tx_mode is set to QUIET. Furthermore, the transmitter's peak-to-peak differential output voltage shall be greater than 720 mV within 500 ns of tx_mode being set to ALERT. The transmitter output shall be fully compliant within 5 μs after tx_mode is set to DATA. During LPI mode, the common-mode voltage shall be maintained to within ± 150 mV of the pre-LPI value.

128.7.1.5 Differential output return loss

For frequencies from 100 MHz to 2000 MHz, the differential output return loss, in dB, of the transmitter shall meet the requirements of Table 128-4. This output impedance requirement applies to all valid output levels. The reference impedance for differential output return loss measurements shall be 100 Ω.

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 10 & 100 \leq f < 625 \\ 10 - 10\log_{10}(f/625) & 625 \leq f \leq 2000 \end{array} \right\} \text{(dB)} \tag{128-3}$$

where

f is the frequency in MHz

The minimum differential output return loss is shown in Figure 128-4.

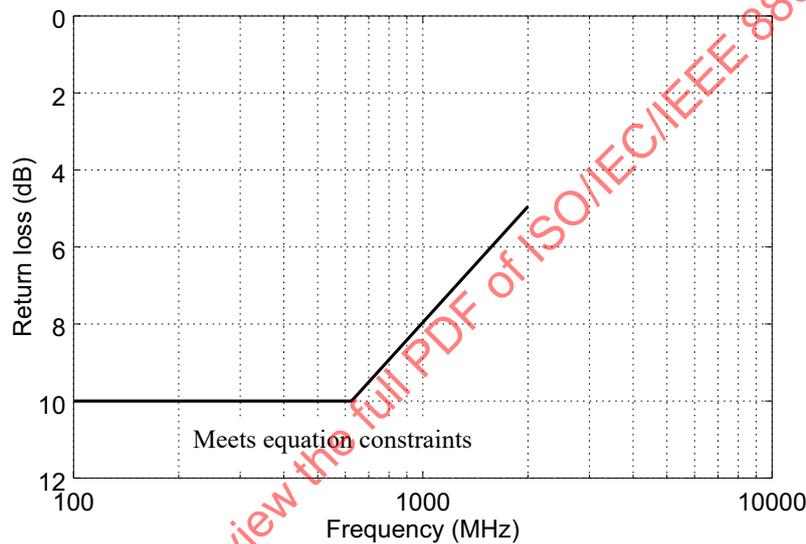


Figure 128-4—Transmitter differential-mode return loss limit

128.7.1.6 Common-mode output return loss

The transmitter common-mode return loss shall meet the requirements of Table 128-4. The reference impedance for common-mode return loss measurements is 25 Ω.

$$Return_loss(f) \geq \left\{ \begin{array}{ll} 7 & 100 \leq f < 625 \\ 7 - 10\log_{10}(f/625) & 625 \leq f \leq 2000 \end{array} \right\} \text{(dB)} \tag{128-4}$$

where

f is the frequency in MHz

The minimum common-mode output return loss is shown in Figure 128-5.

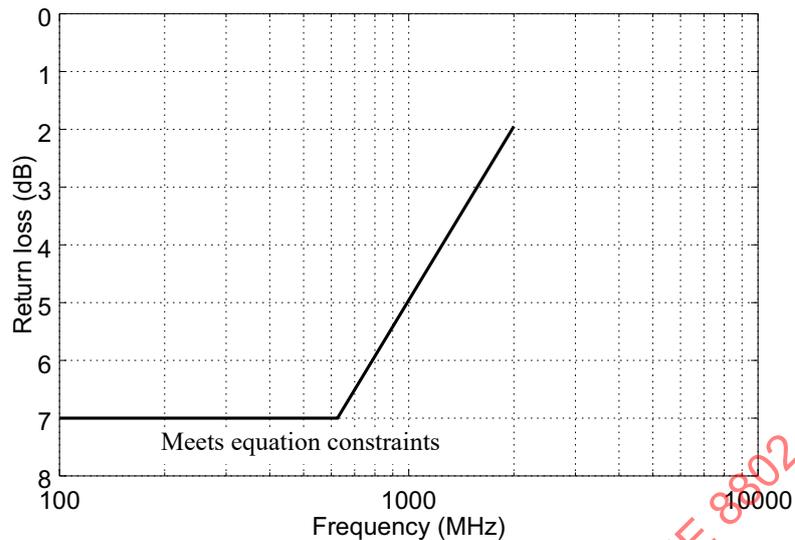


Figure 128-5—Transmitter common-mode return loss limit

128.7.1.7 Transition time

The rising and falling edge transition times shall be as specified in Table 128-4 when measured with respect to the peak-to-peak differential value of the waveform using the high-frequency test pattern of 48A.1.

128.7.1.8 Transmit jitter test requirements

Transmit jitter is defined with respect to a test procedure resulting in a BER bathtub curve such as that described in Annex 48B.3. For the purpose of jitter measurement, the effect of a single-pole high-pass filter with a 3 dB point at 1.875 MHz is applied to the jitter. The data pattern for jitter measurements shall be a high-frequency test pattern as defined in 48A.1. Crossing times are defined with respect to the mid-point (0 V) of the AC-coupled differential signal.

The duty cycle distortion test pattern shall be the high-frequency test pattern as defined in 48A.1.

128.7.1.9 Transmit jitter

The transmit jitter shall be as specified in Table 128-4. Duty cycle distortion (DCD) is considered a component of deterministic jitter. The peak-to-peak duty cycle distortion is defined as the absolute value of the difference in the mean pulse width of a 1 pulse or the mean pulse width of a 0 pulse (as measured at the mean of the high- and low-voltage levels in a clock like repeating 0101 bit sequence) and the nominal pulse width. Jitter specifications are specified for BER = 10^{-12} . Transmit jitter test requirements are specified in 128.7.1.8.

NOTE—Duty cycle distortion (DCD) is also referred to as Even-odd jitter (see 92.8.3.8.1).

128.7.2 Receiver characteristics

Receiver characteristics at TP4 (see Figure 129–1) are summarized in Table 128–5 and detailed in 128.7.2.1 through 128.7.2.5.

Table 128–5—Receiver characteristics for 2.5GBASE-KX

Parameter	Subclause reference	Value	Units
Bit error ratio	128.7.2.1	10^{-12}	
Signaling speed	128.7.2.2	3.125 ± 100 ppm	GBd
Differential input peak-to-peak amplitude (max) ^a	128.7.2.4	1200	mV
Differential input return loss (min) ^b	128.7.2.5	See Equation (128–3)	dB

^aThe receiver shall tolerate amplitudes up to a maximum voltage of 1200mV without permanent damage.

^bRelative to 100 Ω differential.

128.7.2.1 Receiver interference tolerance

The receiver interference tolerance consists of the test as described in Annex 69A with the parameters specified in Table 128–6. The data pattern for the interference tolerance test shall be the test pattern as defined in 48A.5. The receiver shall satisfy the requirements for interference tolerance specified in Annex 69A for the test.

Table 128–6—2.5GBASE-KX interference tolerance parameters

Parameter	Value	Units
Target BER	10^{-12}	
m_{TC} ^a (min)	1	
Amplitude of broadband noise (min RMS)	8.1	mV
Applied transition time (20%–80%, min)	100	ps
Applied sinusoidal jitter (min peak-to-peak)	0.12	UI
Applied random jitter (min peak-to-peak) ^b	0.2	UI

^a m_{TC} is defined in Equation 69A-6 in Annex 69A.

^bApplied random jitter is specified at a BER of 10^{-12} .

128.7.2.2 Signaling speed range

The signaling speed of a 2.5GBASE-KX receiver shall comply with the requirements of Table 128–5.

128.7.2.3 AC-coupling

The 2.5GBASE-KX receiver shall be AC-coupled to the backplane to allow for maximum interoperability between various 2.5 Gb/s components. AC-coupling is considered to be part of the receiver for the purposes of this specification unless explicitly stated otherwise. It should be noted that there may be various methods for AC-coupling in actual implementations.

NOTE—It is recommended that the maximum value of the coupling capacitors be limited to 100 nF. This will limit the inrush currents to the receiver that could damage the receiver circuits when repeatedly connected to transmit modules with a higher voltage level.

128.7.2.4 Input signal amplitude

2.5GBASE-KX receivers shall accept differential input signal peak-to-peak amplitudes produced by compliant transmitters connected without attenuation to the receiver, and still meet the BER requirement specified in 128.7.2.1. Note that this may be larger than the differential maximum of 128.7.1.4 due to the actual transmitter output and receiver input impedances. The input impedance of a receiver can cause the minimum signal into a receiver to differ from that measured when the receiver is replaced with a 100 Ω test load. Since the receiver is AC-coupled, the absolute voltage levels with respect to the receiver ground are dependent on the receiver implementation.

128.7.2.5 Differential input return loss

For frequencies from 100 MHz to 2000 MHz, the differential input return loss, in dB, of the receiver shall meet the requirements of Table 128–5.

128.8 Interconnect characteristics

Informative interconnect characteristics are provided in Annex 69B.

128.9 Environmental specifications

128.9.1 General safety

All equipment that meets the requirements of this standard shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

128.9.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

128.9.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

128.9.4 Electromagnetic compatibility

A system integrating the 2.5GBASE-KX PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

128.9.5 Temperature and humidity

A system integrating the 2.5GBASE-KX PHY is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

128.10 Protocol implementation conformance statement (PICS) proforma for Clause 128, Physical Medium Dependent sublayer and baseband medium, type 2.5GBASE-KX⁶

128.10.1 Introduction

The supplier of a protocol implementation that is claimed to conform to IEEE Std 802.3, Clause 128, Physical Medium Dependent sublayer and baseband medium type 2.5GBASE-KX, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

128.10.2 Identification

128.10.2.1 Implementation identification

Supplier	
Contact point for inquiries about the PICS	
Implementation Name(s) and Version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)	
NOTE 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for the identification. NOTE 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

128.10.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3cb-2018, Clause 128, Physical Medium Dependent sublayer and baseband medium type 2.5GBASE-KX
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3cb-2018.)	
Date of Statement	

⁶Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

128.10.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGMII	XGMII	128.1, Clause 46	Interface is supported	O	Yes [] No []
PCS	Support of 2.5GBASE-X PCS/PMA	128.1, 73.9		M	Yes []
AN	Auto-Negotiation for Back- plane Ethernet	128.1, Clause 73	Device implements Auto- Negotiation for Backplane Ethernet	M	Yes []
PCS	PCS associated with PMD sup- ports AN service primitive AN_LINK	128.3, 73.9	Required generation of PMD primitive	M	Yes []
DC	Delay Constraints	128.4	Device conforms to delay constraints	M	Yes []
*MD	MDIO Capability	128.5	If MDIO is implemented, map control variables as described in 128.5	O	Yes [] No []
*LPI	LPI	128.6.10	LPI function implemented if EEE is supported	O	Yes [] No []
*SD	Signal Detect Generation	128.6.4	Signal detect implemented if EEE is supported	LPI:M	Yes [] N/A []
*TD	PMD_transmit_disable	128.6.5	PMD transmit disable function implemented if EEE is supported	LPI:M	Yes [] N/A []

128.10.4 PICS proforma tables for Clause 128, Physical Medium Dependent (PMD) sublayer and baseband medium, type 2.5GBASE-KX.

128.10.4.1 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	PMD_global_signal_detect set to the value of SIGNAL_DETECT	128.2.3.2	Required generation of PMD primitive	M	Yes []
FS2	Tx+Rx delays from 2.5GBASE-KX PMD and medium ≤ 256 bit times	128.4	Required timing	M	Yes []
FS3	PMD Transmit function	128.6.2	Conveys bits from PMD service interface to MDI	M	Yes []
FS4	PMD Transmitter signal	128.6.2	A positive differential voltage corresponds to tx_bit = ONE	M	Yes []
FS5	PMD Receive function	128.6.3	Conveys bits from MDI to PMD service interface	M	Yes []
FS6	PMD Receiver signal	128.6.3	A positive differential voltage corresponds to rx_bit = ONE	M	Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
FS7	PMD Signal Detect function - not implemented	128.6.4	Continuously reported OK via PMD_SIGNAL.indication (SIGNAL_DETECT).	!SD:M	Yes [] N/A []
FS8	PMD Signal Detect during LPI - implemented	128.6.4	Report to the PMD service interface using the message PMD_SIGNAL.indication (SIGNAL_DETECT) if EEE supported	LPI:M	Yes [] N/A []
FS9	PMD Signal Detect timing LPI quiet to active	128.6.4	Set SIGNAL_DETECT to OK within 750 ns	LPI:M	Yes [] N/A []
FS10	PMD Signal Detect timing active to LPI quiet	128.6.4	Set SIGNAL_DETECT to FAIL within 750 ns	LPI:M	Yes [] N/A []
FS11	Transmit Disable	128.6.5	Disables Transmitter when PMD_Transmit_disable set to ONE	TD:M	Yes [] N/A []
FS12	Transmit Disable during LPI	128.6.5	Disable transmitter during tx_quiet	LPI:M	Yes [] N/A []
FS13	PMD_fault	128.6.5	Transmit disabled if detected	TD:O	Yes [] No [] N/A []
FS14	PMD_transmit_disable	128.6.5	Loopback function not affected	TD:M	Yes [] N/A []
FS15	tx_quiet disabled transmitter	128.7.1	Disables Transmitter when tx_quiet is asserted as specified in 128.7.1.5	LPI:M	Yes [] N/A []
FS16	Loopback Function	128.6.6	Loopback function provided	M	Yes []
FS17	Loopback effect on Transmitter	128.6.6	Loopback function does not disable transmitter	M	Yes []

128.10.4.2 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	MDIO Variable Mapping	128.5	Map MDIO variables to PMD variables as shown in Table 128-2 and Table 128-3	MD:M	Yes [] N/A []
MF2	PMD_fault function	128.6.7	Sets PMD_fault to a logical 1 if any local fault is detected; otherwise, set to 0	MD:M	Yes [] N/A []
MF3	PMD_transmit_fault function	128.6.8	Sets PMD_transmit_fault to a logical 1 if any local fault is detected on the transmit path; otherwise, set to 0	MD:M	Yes [] N/A []
MF4	PMD_receive_fault function	128.6.9	Sets PMD_receive_fault to a logical 1 if any local fault is detected on the receive path; otherwise, set to 0	MD:M	Yes [] N/A []

128.10.4.3 Transmitter electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	100 Ω differential test fixture	128.7.1.2	As shown in Figure 128–2	M	Yes []
TC2	Differential return loss of test fixture	128.7.1.2	Per Equation (128–1) and Equation (128–2)	M	Yes []
TC3	Signaling speed	128.7.1.3	3.125 GBd ± 100 ppm	M	Yes []
TC4	Maximum transmitter differential peak-to-peak voltage	128.7.1.4	1200 mV, pk-pk, using the pattern as defined in 48A.2	M	Yes []
TC5	Minimum transmitter differential peak-to-peak voltage	128.7.1.4	800 mV, pk-pk, using the pattern as defined in 48A.2	M	Yes []
TC6	Maximum transmitter differential peak-to-peak voltage when TX disabled	128.6.5	30 mV, pk-pk	M	Yes []
TC7	Common-mode output voltage	128.7.1.4	0V to +1.9V	M	Yes []
TC8	Common-mode output return loss	128.7.1.6	Per Equation (128–3)	M	Yes []
TC9	Reference Impedance	128.7.1.6	100 Ω for differential return loss measurements	M	Yes []
TC10	Transmit jitter, peak-to-peak	128.7.1.8	Max TJ of 0.32 UI. Max DJ of 0.12 UI	M	Yes []
TC11	Output Amplitude LPI voltage	128.7.1.4	Less than 30 mV within 500 ns of tx_mode = QUIET	LPI:M	Yes [] N/A []
TC12	Output Amplitude ON voltage	128.7.1.4	Greater than 720 mV within 500 ns of tx_mode = ALERT	LPI:M	Yes [] N/A []
TC13	Output to DATA compliance timing	128.7.1.4	Fully compliant ≤ 5 μsec from tx_mode = DATA	LPI:M	Yes [] N/A []
TC14	Common-mode voltage in LPI mode	128.7.1.4	output within ±150 mV of the pre-LPI value	LPI:M	Yes [] N/A []
TC15	Tx output waveform measurement	128.7.1.4	Measured as shown in Figure 128–3	M	Yes []
TC16	Differential Output Return Loss	128.7.1.5	Equation (128–3) and Equation (128–3) at reference impedance of 100 Ω	M	Yes []
TC17	TX common-mode return loss	48A.1	Equation (128–4)	M	Yes []
TC18	TX rise time and fall time	128.7.1.7	30 ps to 100 ps, 20%/80% of peak-to-peak differential with pattern in 48A.1	M	Yes []
TC19	Jitter test pattern	128.7.1.8	See pattern definition in 128.7.1.8	M	Yes []
TC20	DCD test pattern	128.7.1.8	See pattern definition in 128.7.1.8	M	Yes []
TC21	TX jitter limit	128.7.1.9	≤ 0.32 UI, peak-to-peak	M	Yes []
TC22	TX DCD limit	128.7.1.9	≤ 0.035 UI, peak-to-peak	M	Yes []

128.10.4.4 Receiver electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver interference tolerance measurement method	128.7.2.1	Per Annex 69A with parameters specified in Table 128–6	M	Yes []
RC2	Receiver interference tolerance test pattern	128.7.2.1	test patterns 2 or 3 as defined in 48A.4	M	Yes []
RC3	Receiver interference tolerance requirements	128.7.2.1	Satisfy requirements per Annex 69A	M	Yes []
RC4	Input signaling speed in the range of 3.125 GBd ± 100 ppm	128.7.2.2	Receiver meets requirements of Table 128–5	M	Yes []
RC5	Receiver coupling	128.7.2.3	AC-coupled	M	Yes []
RC6	Input signal amplitude	128.7.2.4	BER still met when compliant transmitter is connected with no attenuation	M	Yes []
RC7	Differential input return loss	128.7.2.5	Per Equation (128–3)	M	Yes []
RC8	Reference Impedance	128.7.2.5	100 Ω for differential return loss measurements	M	Yes []

128.10.4.5 Environmental and safety specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	128.9.1	Conforms to IEC 60950-1	M	Yes []
ES2	Electromagnetic compatibility	128.9.4	Comply with applicable local and national codes	M	Yes []

129. Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R

129.1 Overview

129.1.1 Scope

This clause specifies the Physical Coding Sublayer (PCS) and the Physical Medium Attachment (PMA) Sublayer that is common to a family of 5 Gb/s Physical Layer implementations, known as 5GBASE-R. The 5GBASE-R PCS and 5GBASE-R PMA are sublayers of the 5 Gb/s BASE-R PHY listed in Table 125-1. The term 5GBASE-R is used when referring generally to Physical Layers using the PCS and PMA defined in this clause.

129.1.2 Relationship of 5GBASE-R to other standards

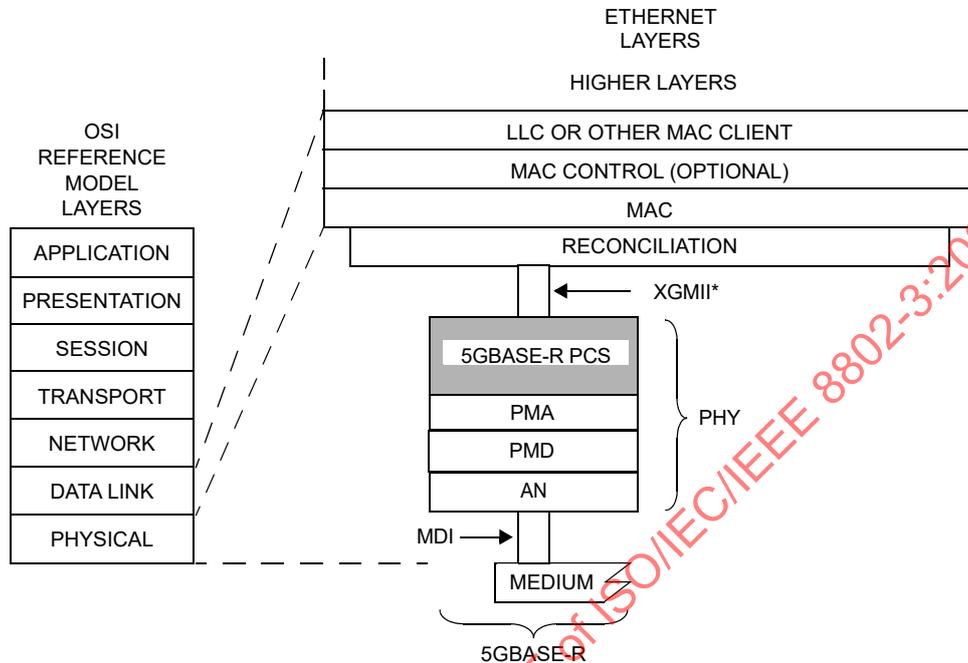
Figure 125-1 depicts the relationships among the 5GBASE-R sublayers, the Ethernet MAC and reconciliation layers, and the higher layers. The 5GBASE-R service interface is the XGMII, which is defined in Clause 46.

The 5GBASE-R PCS is identical to the 10GBASE-R PCS specified in Clause 49 with the exception that the BER Monitor process asserts hi_ber if ber_cnt reaches 16 in a 250 μ s period. This differs from the definition in 49.2.13.3 where it asserts hi_ber if ber_cnt reaches 16 in a 125 μ s period.

The purpose of the serial PMA is to attach the PMD of choice to its client. The PMA service interface is defined in an abstract manner and does not imply any particular implementation. The nominal bit rate of the PMA is 5.15625 Gb/s.

129.1.3 Summary of 5GBASE-R sublayers

Figure 129–1 depicts the relationship between the 5GBASE-R PCS and its associated sublayers.



AN = AUTO-NEGOTIATION
 MAC = MEDIUM ACCESS CONTROL
 LLC = LOGICAL LINK CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT

PMD = PHYSICAL MEDIUM DEPENDENT
 XGMII = 10 GIGABIT MEDIA INDEPENDENT INTERFACE
 * XGMII IS OPTIONAL

Figure 129–1—5GBASE-R PCS and PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

129.1.3.1 Physical Coding Sublayer (PCS)

The PCS service interface is the XGMII, which is defined in Clause 46 operating at 5 Gb/s. The 5GBASE-R PCS provides all services required by the XGMII, including encoding (decoding) of eight XGMII data octets to (from) 66-bit blocks (64B/66B) for communication with the underlying PMA.

129.1.3.2 Physical Medium Attachment (PMA) sublayer

The PMA provides a medium-independent means for the PCS to support the use of serial-bit oriented physical media. The 5GBASE-R PMA performs the following functions:

In the transmit direction (i.e., transmitting data from the PMA client to the PMD), the PMA performs the following functions:

- a) Provide transmit source clock to PMA client
- b) Serialization of 16-bit data to serial bit stream.
- c) Transmission of serial data to PMD.

In the receive direction (i.e., serial data from PMD to the PMA client), the PMA performs the following functions:

- a1) Bit clock recovery of serial data from PMD.
- b1) Provide receive clock to PMA client.
- c1) Deserialization of serial data to 16-bit parallel data.
- d1) Transmission of parallel data to PMA client.
- e1) Provide link status information.

129.1.4 Inter-sublayer interfaces

There are a number of interfaces employed by 5GBASE-R. Some (such as the PMA service interface) use an abstract service model to define the operation of the interface. The PCS service interface is the XGMII running at 5 Gb/s that is defined in Clause 46. The XGMII has an optional physical instantiation.

The upper interface of the PCS may connect to the Reconciliation Sublayer through the XGMII. The lower interface of the PCS connects to the PMA sublayer to support a 5GBASE-KR PMD. The nominal rate of the PMA service interface is 322.265625 Mtransfers/s, which provides capacity for the MAC data rate of 5 Gb/s.

The nominal rate of the PMD service interface is 5.15625 Gb/s.

If the optional Energy-Efficient Ethernet (EEE) capability is supported (see Clause 78, 78.3), then the interface with the PMA sublayer includes rx_mode and tx_mode to control power states in lower sublayers and energy_detect that indicates whether the PMD sublayer has detected a signal at the receiver.

It is important to note that, while this specification defines interfaces in terms of bits, octets, and frames, implementations may choose other data-path widths for implementation convenience.

129.2 Physical Coding Sublayer (PCS)

129.2.1 Functions within the PCS

The 5GBASE-R PCS shall have all the functionality of the 10GBASE-R PCS specified in Clause 49.

The BER monitor state diagram shown in Figure 49–15 applies but the definition of “ber_cnt” in 49.2.13.2.4 is replaced with “Count up to a maximum of 16 of the number of invalid sync headers within the current 250 μ s period.” and the ber_cnt definition in 49.2.14.2 is similarly modified.

The PCS encodes data and control information into 66-bit blocks. The relationship of block bit positions to the XGMII, PMA, and other PCS constructs is illustrated in Figure 129–2 for transmit and Figure 129–3 for receive (see figures in 129.2.3).

129.2.2 Notation conventions

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

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

129.2.3 Transmission order

Block bit transmission order is illustrated in Figure 129–2 and Figure 129–3. These figures show the mapping from XGMII to 64B/66B block for a block containing eight data characters.

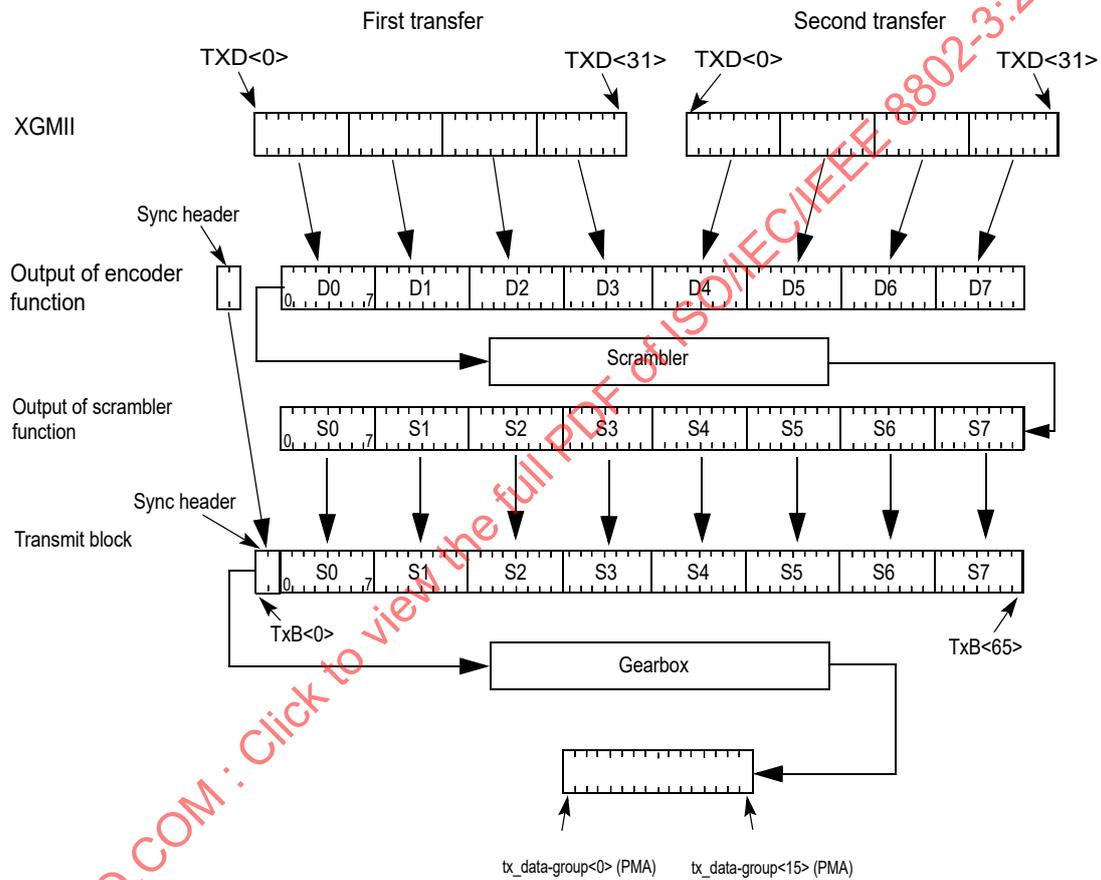


Figure 129–2—PCS Transmit bit ordering

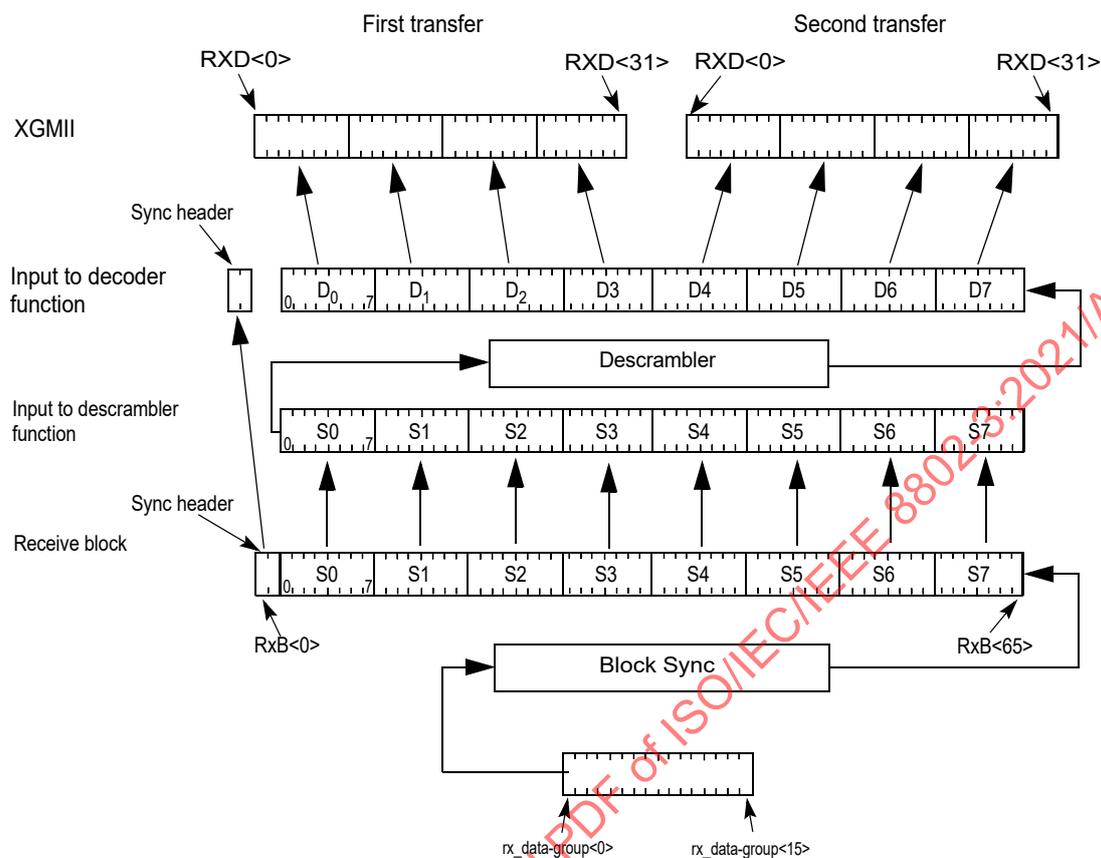


Figure 129-3—PCS Receive bit ordering

129.2.4 Low Power Idle

If the 5GBASE-KR PCS is a part of a PHY configured for EEE operation, the PCS shall follow the state diagrams specified in Figure 49-12 and Figure 49-13.

The LPI functions shall use timer values for these state diagrams as shown in Table 129-1 for transmit and Table 129-2 for receive.

Table 129-1—Transmitter LPI timing parameters

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

Table 129–2—Receiver LPI timing parameters

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

129.2.5 PCS used with 5GBASE-KR PMD

The following requirements apply to a PCS used with a 5GBASE-KR PMD. Support for the Auto-Negotiation process defined in Clause 73 is mandatory. The PCS shall support the primitive AN_LINK.indication(link_status) (see 73.9). The parameter link_status shall take the value FAIL when PCS_status=false and the value OK when PCS_status=true. The primitive shall be generated when the value of link_status changes.

129.3 Physical Medium Attachment (PMA) sublayer

129.3.1 Service Interface

The Serial PMA provides a Service Interface to the 5G BASE-R PCS. These services are described in an abstract manner and do not imply any particular implementation. The PMA Service Interface shall support the exchange of data-groups between the PMA and the PMA client. The PMA converts data-groups into bits and passes these to the PMD, and vice versa. It also generates an additional status indication for use by its client.

The following primitives are defined:

- PMA_UNITDATA.request(tx_data-group<15:0>)
- PMA_UNITDATA.indication(rx_data-group<15:0>)
- PMA_SIGNAL.indication(SIGNAL_OK)
- PMA_RXMODE.request(rx_mode)
- PMA_TXMODE.request(tx_mode)
- PMA_ENERGY.indication(energy_detect)

The definitions of these primitives are found in 51.2.

129.3.2 Functions within the PMA

The PMA comprises the PMA Transmit and PMA Receive processes for 5GBASE-R. The PMA Transmit process serializes the tx_data-groups and passes them to the PMD for transmission on the underlying medium. Similarly, the PMA Receive process deserializes received data from the PMD and presents the data as rx_data-groups to the PMA client. The PMA receiver continuously conveys 16-bit data-groups to the PMA client, independent of data-group alignment.

129.3.2.1 PMA transmit function

The PMA Transmit function passes data unaltered (except for serializing) from the PMA client directly to the PMD. Upon receipt of a PMA_UNITDATA.request primitive, the PMA Transmit function shall serialize the 16 bits of the tx_data-group<15:0> parameter and transmit them to the PMD in the form of 16 successive PMD_UNITDATA.request primitives.

129.3.2.2 PMA receive function

The PMA Receive function passes data unaltered from the PMD directly to the PMA client. Upon receipt of 16 successive PMD_UNITDATA.indication primitives, the PMA shall assemble the 16 received bits into a single 16-bit value and pass that value to the PMA client as the rx_data-group<15:0> parameter of the primitive PMA_UNITDATA.indication. The PMA receive function does not align rx_datagroup<15:0> to the original tx_data-group<15:0> from the remote end of the link.

129.3.3 PMA loopback mode (optional)

PMA loopback is optional. If PMA loopback is implemented, it shall conform to the requirements of 51.8.

129.4 Compatibility considerations

There is no requirement for a compliant device to implement or expose any of the interfaces specified for the PCS or PMA. Implementations of a XGMII shall comply with the requirements as specified in Clause 46.

129.5 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must conform to certain delay maxima, and that network planners and administrators conform to constraints regarding the cable topology and concatenation of devices.

The sum of transmit and receive delay contributed by the 5GBASE-R PCS and PMA shall be no more than 3584 bit-times.

129.6 Environmental specifications

All equipment subject to this clause shall conform to the requirements of 51.9.

129.7 Protocol implementation conformance statement (PICS) proforma for Clause 129, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R⁷

129.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 129, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

129.7.2 Identification

129.7.2.1 Implementation identification

Supplier ¹	
Contact point for inquiries about the PICS ¹	
Implementation Name(s) and Version(s) ^{1,3}	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s) ²	
NOTE 1—Required for all implementations. NOTE 2—May be completed as appropriate in meeting the requirements for the identification. NOTE 3—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

129.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3cb-2018, Clause 129, Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) sublayer for 5 Gb/s 64B/66B, type 5GBASE-R.
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No <input type="checkbox"/> Yes <input type="checkbox"/> (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3cb-2018.)	
Date of Statement	

⁷Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

129.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGE	XGMII compatibility interface	129.1.4	Compatibility interface is supported	O	Yes [] No []
MD	MDIO	Clause 45, 49.2.14	Registers and interface supported	O	Yes [] No []
PCS	5GBASE-R PCS	129.2.1, 49	Support PCS in Clause 49 as modified in 129.2.1	M	Yes []
BER	BER Monitor test counter	129.2.1, 49	ber_cnt of 16 in a period of 250 μs	M	Yes []
*LPI	Implementation of LPI	Clause 78		O	Yes [] No []

129.7.4 PICS Proforma Tables for PCS, type 5GBASE-R

129.7.4.1 Coding rules

Item	Feature	Subclause	Value/Comment	Status	Support
C1	Encoder implements the code as specified	49.2.13.2.3 and 49.2.4		M	Yes []
C2	Decoder implements the code as specified	49.2.13.2.3 and 49.2.4		M	Yes []
C3	IDLE control code insertion and deletion	49.2.4.7	Insertion or Deletion in groups of 4 /I/s	M	Yes []
C4	IDLE control code deletion	49.2.4.7	When deleting /I/s, the first four characters after a /T/ shall not be deleted.	M	Yes []
C5	Sequence ordered set deletion	49.2.4.10	Only one whole ordered set of two consecutive sequence ordered sets may be deleted	M	Yes []

129.7.4.2 Scrambler and Descrambler

Item	Feature	Subclause	Value/Comment	Status	Support
S1	Scrambler	49.2.6	Performs as shown in Figure 49–8	M	Yes []
S2	Descrambler	49.2.10	Performs as shown in Figure 49–10	M	Yes []

129.7.5 Test-pattern modes

Item	Feature	Subclause	Value/Comment	Status	Support
JT1	Square wave and pseudo-random transmit test-pattern generators are implemented	49.2.8	Performs as in 49.2.8	M	Yes []
JT2	Pseudo-random receive test-pattern checker is implemented	49.2.12	Performs as in 49.2.12	M	Yes []
*JT3	Support for PRBS31 test pattern	49.2.8		O	Yes [] No []
JT4	PRBS31 test-pattern generator is implemented	49.2.8	Performs as in 49.2.8	JT3:M	Yes [] N/A []
JT5	PRBS31 test-pattern checker is implemented	49.2.12	Performs as in 49.2.12	JT3:M	Yes [] N/A []
*JT6	Support for PRBS9 transmit test pattern	49.2.8		O	Yes [] No []
JT7	PRBS9 transmit test pattern is implemented	49.2.8	Performs as in 49.2.8	JT6:M	Yes [] N/A []

129.7.5.1 Bit order

Item	Feature	Subclause	Value/Comment	Status	Support
B1	Transmit bit order	49.2.7	Placement of bits into tx_data-group<15:0> as specified in 49.2.7	M	Yes []
B2	Receive bit order	49.2.9	Placement of bits from rx_data-group<15:0> into blocks as specified in 49.2.9	M	Yes []

129.7.6 Management

Item	Feature	Subclause	Value/Comment	Status	Support
M1	Alternate access to PCS Management objects is provided	49.2.14		O	Yes [] No []

129.7.6.1 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SM1	Lock	49.2.13.3	Meets the requirements of Figure 49-14	M	Yes []
SM2	BER Monitor	49.2.13.3	Meets the requirements of Figure 49-15	M	Yes []
SM3	Transmit	49.2.13.3	Meets the requirements of Figure 49-16	M	Yes []
SM4	Receive	49.2.13.3	Meets the requirements of Figure 49-17	M	Yes []

129.7.6.2 Loopback

Item	Feature	Subclause	Value/Comment	Status	Support
LB1	PMA Loopback	129.3.3	Conform to the requirements of 51.8	O	Yes [] No []

129.7.6.3 Delay Constraints

Item	Feature	Subclause	Value/Comment	Status	Support
TIM1	PCS Delay Constraint	49.2.15	No more than 3584 BT for sum of transmit and receive path delays	M	Yes []

129.7.6.4 Auto-Negotiation for Backplane Ethernet functions

Item	Feature	Subclause	Value/Comment	Status	Support
AN1	Support for use with a 5GBASE-KR PMD	49.2.16	AN technology dependent interface described in Clause 73	M	Yes []
AN2	AN_LINK.indication primitive	49.2.16 and 129.2.5	Support of the primitive AN_LINK.indication(link_status), when the PCS is used with 5GBASE-KR PMD. Primitive generated when link_status changes.	M	Yes []
AN3	link_status parameter	49.2.16 and 129.2.5	Takes the value OK or FAIL, as described in 49.2.16	M	Yes []
AN4	Generation of AN_LINK.indication primitive	49.2.16 and 129.2.5	Generated when the value of link_status changes	M	Yes []

129.7.6.5 LPI functions

Item	Feature	Subclause	Value/Comment	Status	Support
LP1	Insertion and deletion of LPIs in groups of four	49.2.4.7		LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP2	Unscrambled data transmitted when scrambler_bypass = TRUE	49.2.6		LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP3	Scrambler continues to operate as normal when scrambler_bypass = TRUE	49.2.6		LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP4	Transmit state diagrams	49.2.13.3	Support additions to Figure 49–16 for LPI operation	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP5	Receive state diagrams	49.2.13.3	Support additions to for LPI operation	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP6	LPI transmit state diagrams	49.2.13.3.1	Meets the requirements of Figure 49–12	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP7	LPI receive state diagrams	49.2.13.3.1	Meets the requirements of Figure 49–13	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP8	LPI transmit timing	129.2.4, 49.2.13.3.1	Meets the requirements of Table 129–1 instead of Table 49–2	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]
LP9	LPI receive timing	129.2.4, 49.2.13.3.1	Meets the requirements of Table 129–2 instead of Table 49–3	LPI:M	Yes [<input type="checkbox"/>] N/A [<input type="checkbox"/>]

130. Physical Medium Dependent sublayer and baseband medium, type 5GBASE-KR

130.1 Overview

This clause specifies the 5GBASE-KR PMD and the baseband medium. When forming a complete PHY, a PMD shall be connected to the appropriate sublayers (see Table 130–1), and with the management functions that are optionally accessible through the management interface defined in Clause 45, or equivalent. References to the MDI (Media Dependent Interface) should be considered to be TP1 for the transmitter and TP4 for the receiver, as measurement points.

Table 130–1—Physical Layer clauses associated with the 5GBASE-KR PMD

Associated clause	5GBASE-KR
46—XGMII ^a	Optional
73—Auto Negotiation for Backplane	Required
78—EEE	Optional
129—PCS and PMA for 64B/66B	Required

^aThe XGMII is an optional interface. However, if the XGMII is not implemented, a conforming implementation must behave functionally as though the RS and XGMII were present.

A 5GBASE-KR PHY with the optional Energy-Efficient Ethernet (EEE) capability may optionally enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization.

130.2 Physical Medium Dependent (PMD) service interface

The following specifies the services provided by the 5GBASE-KR PMDs. These PMD sublayer service interfaces are described in an abstract manner and do not imply any particular implementation.

The PMD Service Interface supports the exchange of encoded and scrambled 64B/66B blocks between the PMA and PMD entities. The PMD translates the serialized data of the PMA to and from signals suitable for the specified medium.

The following primitives are defined:

```

PMD_UNITDATA.request(tx_bit)
PMD_UNITDATA.indication(rx_bit)
PMD_SIGNAL.indication(SIGNAL_DETECT)

```

If EEE is supported, the following primitives are also defined on the PMD Service Interface:

```

PMD_RX_MODE.request(rx_mode)
PMD_TX_MODE.request(tx_mode)

```

These messages affect the PCS variables as described in 49.2.13.2.2.

130.2.1 PMD_UNITDATA.request

This primitive defines the transfer of a serial data stream from the PMA to the PMD.

130.2.1.1 Semantics of the service primitive

PMD_UNITDATA.request(tx_bit)

The data conveyed by PMD_UNITDATA.request is a continuous stream of bits. The tx_bit parameter can take one of two values: ONE or ZERO.

130.2.1.2 When generated

The PMA continuously sends the appropriate stream of bits to the PMD for transmission on the medium, at the signaling speed specified in Table 130–4 for 5GBASE-KR PMD types.

130.2.1.3 Effect of receipt

Upon receipt of this primitive, the PMD converts the specified stream of bits into the appropriate signals on the MDI.

130.2.2 PMD_UNITDATA.indication

This primitive defines the transfer of data (in the form of serialized data) from the PMD to the PMA.

130.2.2.1 Semantics of the service primitive

PMD_UNITDATA.indication(rx_bit)

The data conveyed by PMD_UNITDATA.indication is a continuous stream of bits. The rx_bit parameter can take one of two values: ONE or ZERO.

130.2.2.2 When generated

The PMD continuously sends a stream of bits to the PMA corresponding to the signals received from the MDI.

130.2.2.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

130.2.3 PMD_SIGNAL.indication

This primitive is generated by the PMD to indicate the status of the signal being received from the MDI.

130.2.3.1 Semantics of the service primitive

PMD_SIGNAL.indication(SIGNAL_DETECT)

The SIGNAL_DETECT parameter can take on one of two values: OK or FAIL, indicating whether the PMD is detecting signal at the receiver (OK) or not (FAIL). When SIGNAL_DETECT = FAIL, PMD_UNITDATA.indication(rx_bit) is undefined.

NOTE—SIGNAL_DETECT = OK does not guarantee that PMD_UNITDATA.indication(rx_bit) is known good. It is possible for a poor quality link to provide sufficient signal for a SIGNAL_DETECT = OK indication and still not meet the BER objective.

130.2.3.2 When generated

The PMD generates this primitive to indicate a change in the value of SIGNAL_DETECT. If the MDIO interface is implemented, then PMD_global_signal_detect shall be continuously set to the value of SIGNAL_DETECT.

130.2.3.3 Effect of receipt

The effect of receipt of this primitive by the client is unspecified by the PMD sublayer.

130.2.4 PMD_RX_MODE.request

This primitive is generated by the PCS Receive Process when EEE is supported to indicate that the input signal is quiet and the PMA and PMD receiver may go into a low power mode. When EEE is not supported, the primitive is never invoked, and the PMD behaves as if rx_mode = DATA.

130.2.4.1 Semantics of the service primitive

PMD_RX_MODE.request(rx_mode)

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

130.2.4.2 When generated

The PCS generates this primitive to request the appropriate PMD receive LPI state.

130.2.4.3 Effect of receipt

When rx_mode is QUIET, the PMD receive function may deactivate functional blocks to conserve energy. When rx_mode is DATA, the PMD receive function operates normally.

130.2.5 PMD_TX_MODE.request

This primitive is generated by the PCS Transmit Process when EEE is supported to indicate that the PMA and PMD transmit functions may go into a low power mode and to disable the PMD transmitter. See 130.6.5. When EEE is not supported, the primitive is never invoked, and the PMD behaves as if tx_mode = DATA.

130.2.5.1 Semantics of the service primitive

PMD_TX_MODE.request(tx_mode)

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

130.2.5.2 When generated

The PCS generates this primitive to request appropriate PMD transmit LPI state.

130.2.5.3 Effect of receipt

When tx_mode is QUIET, the PMD Transmit function may deactivate functional blocks to conserve energy. When tx_mode is ALERT, the PMD Transmit function transmits the alert pattern. When tx_mode is DATA, the PMD Transmit function operates normally.

130.3 PCS requirements for Auto-Negotiation (AN) service interface

The PCS associated with this PMD shall support the AN service interface primitive AN_LINK.indication defined in 73.9.

130.4 Delay constraints

Predictable operation of the MAC Control PAUSE operation (Clause 31, Annex 31B) demands that there be an upper bound on the propagation delays through the network. This implies that MAC, MAC Control sublayer, and PHY implementers must consider the delay maxima, and that network planners and administrators consider the delay constraints regarding concatenation of devices.

The sum of the transmit and the receive delays contributed by the 5GBASE-KR PMD and medium shall be no more than 1024 bit times. It is assumed that the round-trip delay through the medium is 80 bit times.

130.5 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that provide control and status information for and about the PMD. If MDIO is implemented, it shall map MDIO control variables to PMD control variables as shown in Table 130–2 and MDIO status variables to PMD status variables as shown in Table 130–3.

Table 130–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	Control register 1	1.0.15	PMD_reset
Global PMD Transmit Disable	Transmit disable register	1.9.0	Global_PMD_transmit_disable

Table 130–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	Status register 1	1.1.7	PMD_fault
Transmit fault	Status register 2	1.8.11	PMD_transmit_fault
Receive fault	Status register 2	1.8.10	PMD_receive_fault
Global PMD Receive signal detect	Receive signal detect register	1.10.0	Global_PMD_signal_detect

130.6 PMD functional specifications

130.6.1 Link block diagram

For purposes of system conformance, the PMD sublayer is standardized at test points TP1 and TP4 as shown in Figure 130–1. The transmitter and receiver blocks include all off-chip components associated with the respective block. For example, external AC-coupling capacitors, if required, are to be included in the receiver block.

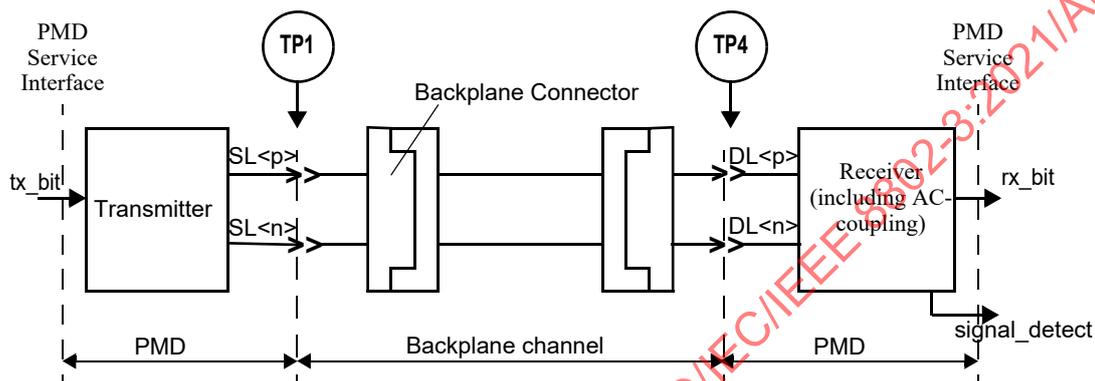


Figure 130–1—Link block diagram

The electrical path from the transmitter block to TP1, and from TP4 to the receiver block, will affect link performance and the measured values of electrical parameters used to verify conformance to this standard. Therefore, it is recommended that this path be carefully designed.

130.6.2 PMD transmit function

The PMD Transmit function shall convey the bits requested by the PMD service interface message `PMD_UNITDATA.request(tx_bit)` to the MDI according to the specifications in this clause. A positive output voltage of $SL<p>$ minus $SL<n>$ (differential voltage) shall correspond to `tx_bit = ONE`.

If the optional Energy-Efficient Ethernet (EEE) capability is supported (see Clause 78), then when `tx_mode` is set to `ALERT`, the PMD shall transmit a repeating 16-bit pattern, hexadecimal `0xFF00`.

130.6.3 PMD receive function

The PMD Receive function shall convey the bits received from the MDI according to the electrical specifications in this clause to the PMD service interface using the message `PMD_UNITDATA.indication(rx_bit)`. A positive input voltage of $DL<p>$ minus $DL<n>$ (differential voltage) shall correspond to `rx_bit = ONE`.

130.6.4 PMD signal detect function

The `Global_PMD_signal_detect` function reports to the PMD service interface, using the message `PMD_SIGNAL.indication(SIGNAL_DETECT)`, which is signaled continuously. `PMD_SIGNAL.indication` is used by 5GBASE-R PHYs as an indicator of valid signal presence. `SIGNAL_DETECT` shall be set to `FAIL` following system reset.

When the PMD signal detect function is implemented, its definition is beyond the scope of this specification. When PMD signal detect is not implemented, the value of SIGNAL_DETECT shall be set to OK for purposes of management and signaling of the primitive. PMD signal detect is mandatory if EEE is supported.

When the PHY supports the optional EEE capability, PMD_SIGNAL.indication is used to indicate when the ALERT signal is detected, which corresponds to the beginning of a refresh or a wake. If the MDIO interface is implemented, then Global_PMD_signal_detect (1.10.0) shall be continuously set to the value of SIGNAL_DETECT as described in 45.2.1.9.7.

When the PHY supports the EEE capability, SIGNAL_DETECT is set to FAIL following a transition from rx_mode = DATA to rx_mode = QUIET. When rx_mode = QUIET, SIGNAL_DETECT shall be set to OK within 500 ns following the application of a signal at the receiver input that is the output of a channel that satisfies the requirements of all the parameters of both interference tolerance test channels defined in 130.7.2.1 when driven by a square wave pattern with a period of 16 unit intervals and peak-to-peak differential output amplitude of 720 mV. While rx_mode = QUIET, SIGNAL_DETECT changes from FAIL to OK only after a valid ALERT signal is applied to the channel.

SIGNAL_DETECT shall be set to FAIL following system reset.

130.6.5 PMD transmit disable function

The Global_PMD_transmit_disable function is mandatory if EEE is supported and is otherwise optional. When this function is supported, it shall meet the following requirements:

- a) When the Global_PMD_transmit_disable variable is set to ONE, this function shall turn off the transmitter such that it drives a constant level (i.e., no transitions) and does not exceed the maximum peak-to-peak differential output voltage specified in Table 130–4.
- b) If a PMD_fault (130.6.7) is detected, then the PMD may turn off the electrical transmitter.
- c) Loopback, as defined in 130.6.6, shall not be affected by Global_PMD_transmit_disable.
- d) For EEE capability, the PMD_transmit_disable function shall turn off the transmitter after tx_mode is set to QUIET within a time and voltage level specified in 130.7.1.4. The PMD_transmit_disable function shall turn on the transmitter after tx_mode is set to DATA or ALERT within the time and voltage level specified in 130.7.1.4.

If the MDIO interface is implemented, then this function shall map to the Global_PMD_transmit_disable bit as specified in 45.2.1.8.7.

130.6.6 Loopback mode

Loopback mode shall be provided for the 5GBASE-KR PMD by the transmitter and receiver of a device as a test function to the device. When loopback mode is selected, transmission requests passed to the transmitter are shunted directly to the receiver, overriding any signal detected by the receiver on its attached link. This bit does not affect the state of the transmitter. The method of implementing loopback mode within the PMD is not defined by this standard.

Control of the loopback function is specified in 45.2.1.1.5.

NOTE 1—The signal path that is exercised in the loopback mode is implementation specific, but it is recommended that this signal path encompass as much of the circuitry as is practical. The intention of providing this loopback mode of operation is to permit diagnostic or self-test functions to test the transmit and receive data paths using actual data. Other loopback signal paths may also be enabled independently using loopback controls within other devices or sublayers.

NOTE 2—Placing a network port into loopback mode can be disruptive to a network.

130.6.7 PMD_fault function

If the MDIO is implemented, PMD_fault is the logical OR of PMD_receive_fault, PMD_transmit_fault, and any other implementation specific fault.

130.6.8 PMD transmit fault function

The PMD_transmit_fault function is optional. The faults detected by this function are implementation specific, but should not include the assertion of the Global_PMD_transmit_disable function.

If a PMD_transmit_fault (optional) is detected, then the Global_PMD_transmit_disable function should also be asserted.

If the MDIO interface is implemented, then this function shall be mapped to the PMD_transmit_fault bit as specified in 45.2.1.7.4.

130.6.9 PMD receive fault function

The PMD_receive_fault function is optional. The faults detected by this function are implementation specific.

If the MDIO interface is implemented, then this function shall contribute to PMA/PMD receive fault bit as specified in 45.2.1.7.5.

130.6.10 PMD LPI function

The PMD LPI function responds to the transitions between Active, Sleep, Quiet, Refresh, and Wake states via the PMD_TX_MODE and PMD_RX_MODE requests. Implementation of the function is optional. The PMD LPI function responds to the transitions between Active, Sleep, Quiet, Refresh, and Wake states via the PMD_TX_MODE and PMD_RX_MODE requests. Implementation of the function is optional. EEE capabilities and parameters will be advertised during the Backplane Auto-negotiation, as described in 45.2.7.15. The transmitter on the local device will inform the link partner's receiver when to sleep, refresh and wake. The local receiver transitions are controlled by the link partner's transmitter and can change independent of the local transmitter states and transitions.

The “Assert LPI” request at the XGMII is encoded in the transmitted symbols. Detection of LPI signaling in the received symbols is indicated as “Assert LPI” at the XGMII. Upon the detection of “Assert LPI” at the XGMII, an energy-efficient 5GBASE-KR PHY continues transmitting for a predefined period, then ceases transmission and deactivates transmit functions to conserve energy. The PHY periodically transmits during this quiet period to allow the remote PHY to refresh its receiver state (e.g., timing recovery, adaptive filter coefficients) and thereby track long-term variations in the timing of the link or the underlying channel characteristics. If, during the quiet or refresh periods, normal interframes resume at the XGMII, the PHY reactivates transmit functions and initiates transmission. This transmission will be detected by the remote PHY, causing it to also exit the LPI mode.

130.7 5GBASE-KR electrical characteristics

130.7.1 Transmitter characteristics

Transmitter characteristics at TP1 (see Figure 130–1) are summarized in Table 130–4 and detailed in 130.7.1.1 through 130.7.1.10.

Table 130–4—Transmitter characteristics for 5GBASE-KR

Parameter	Subclause reference	Value	Units
Signaling speed	130.7.1.3	5.15625 ± 100 ppm	GBd
Peak-to-peak differential output voltage (max)	130.7.1.4	1200	mV
Peak-to-peak differential output voltage (min)	130.7.1.4	800	mV
Peak-to-peak differential output voltage (max) with TX disabled	130.6.5	30	mV
Pre-cursor equalization ratio (R_{pre})	130.7.1.10	1.25 ± 0.05	
Common-mode voltage limits ^c	130.7.1.4	0 to +1.9	V
Common-mode voltage deviation (max) during LPI	130.7.1.4	150	mV
Differential output return loss (min)	130.7.1.5	See Equation (130–3) and Equation (130–4)	dB
Common-mode output return loss (min)	130.7.1.6	See Equation (130–5) and Equation (130–6)	dB
Transition time (20%–80%)	130.7.1.7	20 to 60	ps
Max output jitter (peak-to-peak)			
Random jitter ^a	130.7.1.9	0.15	UI
Deterministic jitter	130.7.1.9	0.12	UI
Duty Cycle Distortion ^b	130.7.1.9	0.035	UI
Total jitter	130.7.1.9	0.27	UI

^aJitter is specified at BER 10^{-12} .

^bDuty Cycle Distortion is considered part of the deterministic jitter distribution.

^cDefined with respect to signal ground as measured at V_{com} in Figure 130–2.

130.7.1.1 Test fixture

The test fixture of Figure 130–2, or its functional equivalent, is required for measuring the transmitter specifications described in 130.7.1, with the exception of return loss.

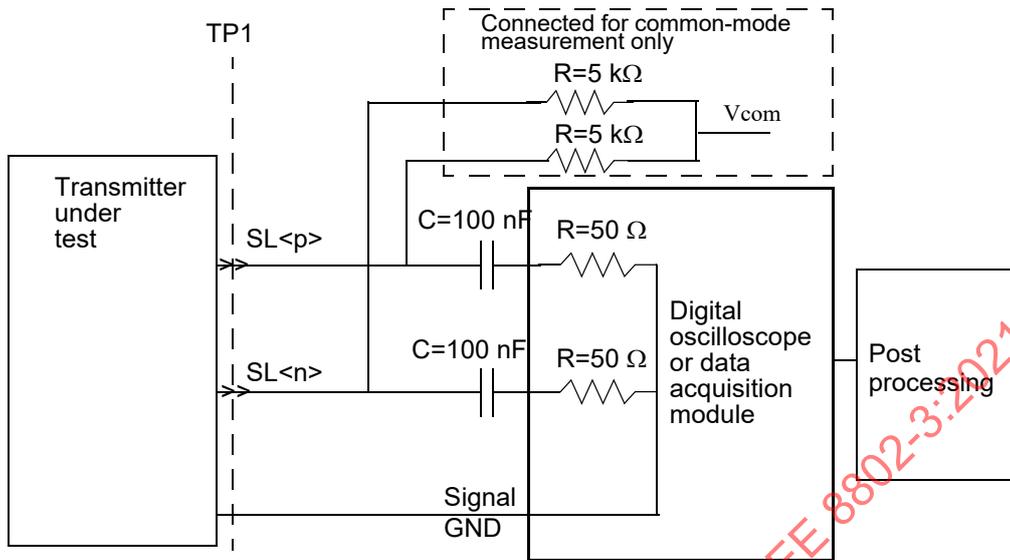


Figure 130-2—Transmit test fixture for 5GBASE-KR

130.7.1.2 Test fixture characteristics

The differential load impedance applied to the transmitter output by the test fixture depicted in Figure 130-2 shall be 100 Ω. The differential return loss, in dB, of the test fixture shall meet the requirements of Equation (130-1) and Equation (130-2).

$$Return_loss(f) \geq 20 \tag{130-1}$$

for 100 MHz ≤ f < 2579 MHz

$$Return_loss(f) \geq Return_loss_{min} = 20 - 26.55 \log_{10} \left(\frac{f}{2579 \text{ MHz}} \right) \tag{130-2}$$

for 2579 MHz ≤ f ≤ 5156.25 MHz

130.7.1.3 Signaling speed

The 5GBASE-KR signaling speed shall be as specified in Table 130-4.

130.7.1.4 Output amplitude

The differential output voltage is constrained via the transmitter output waveform requirements specified in 130.7.1.10. The peak-to-peak differential output voltage shall be as specified in Table 130-4. See Figure 130-3. The differential output voltage test pattern is the square wave test pattern defined in 52.9.1.2, with a run of at least eight consecutive ones followed by an equal number of consecutive zeros.

NOTE—SL<p> and SL<n> are the positive and negative sides of the differential signal pair.

DC-referenced voltage levels are not defined since the receiver is AC-coupled. The common-mode voltage of SL<p> and SL<n>, with respect to signal ground, shall be measured at Vcom in Figure 130-2, with the value as specified in Table 130-4.

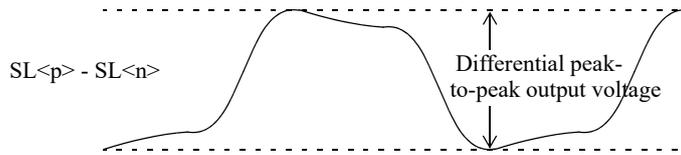


Figure 130-3—Transmitter peak-to-peak differential output voltage definition

For EEE capability, the transmitter shall meet the Table 130-4 requirement for differential peak to peak output voltage when TX is disabled within 500 ns of tx_mode being set to QUIET and remain so while tx_mode is set to QUIET. Furthermore, the transmitter’s peak-to-peak differential output voltage shall be greater than 720 mV within 500 ns of tx_mode being set to ALERT. The transmitter output shall be fully compliant within 5 μs after tx_mode is set to DATA. During LPI mode, the common-mode voltage shall not deviate from the pre-LPI value by more than that allowed in Table 130-4.

130.7.1.5 Differential output return loss

For frequencies from 100 MHz to 3750 MHz, the differential output return loss, in dB, of the transmitter shall meet the requirements of Table 130-4. This output impedance requirement applies to all valid output levels. The reference impedance for differential output return loss measurements shall be 100 Ω.

$$Return_loss \geq 10 \tag{130-3}$$

for 100 MHz ≤ f < 1289 MHz and,

$$Return_loss(f) \geq Return_loss_{min} = 10 - 13.275 \log_{10} \left(\frac{f}{1289 \text{ MHz}} \right) \tag{130-4}$$

for 1289 MHz ≤ f ≤ 3750 MHz.

The minimum differential output return loss is shown in Figure 130-4.

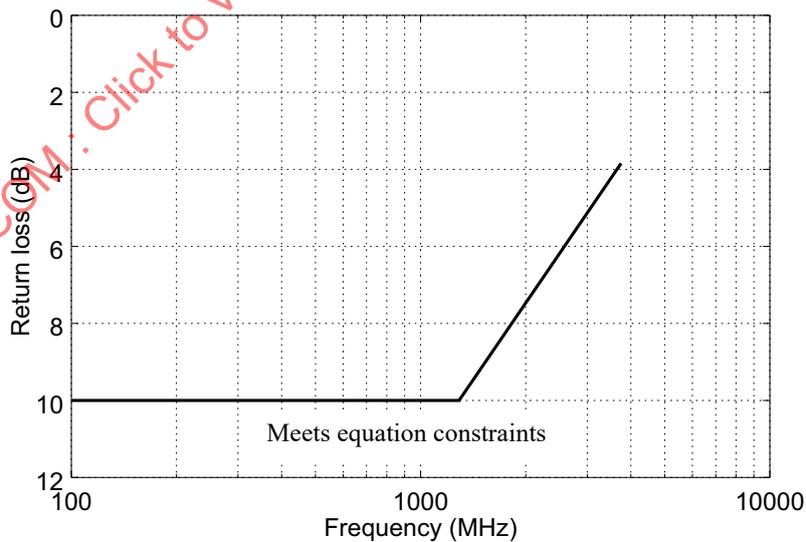


Figure 130-4—Transmit output differential-mode return loss limit

130.7.1.6 Common-mode output return loss

The transmitter common-mode return loss shall meet the requirements of Table 130–4. The reference impedance for common-mode return loss measurements is 25 Ω.

$$Return_loss \geq 7 \tag{130-5}$$

for 100 MHz ≤ f < 1250 MHz,

$$Return_loss(f) \geq 7 - 13.275 \log_{10} \left(\frac{f}{1250 \text{ MHz}} \right) \tag{130-6}$$

for 1250 MHz ≤ f ≤ 3750 MHz.

The minimum common-mode output return loss is shown in Figure 130–5.

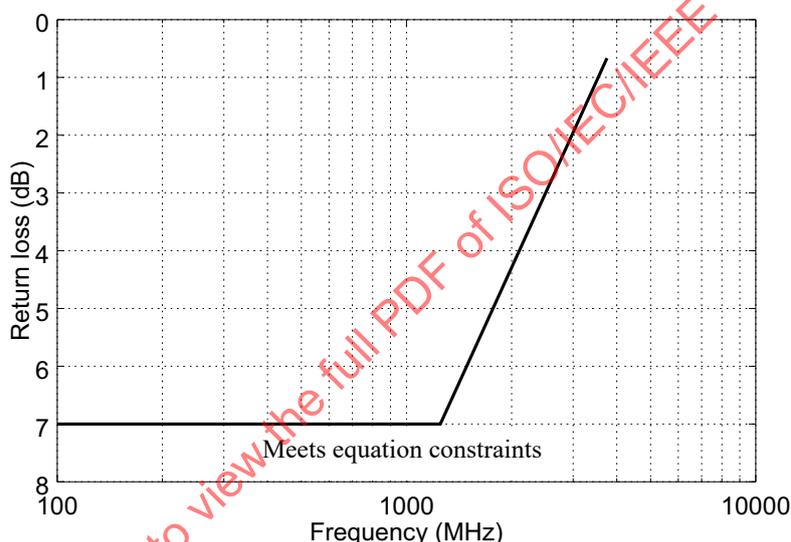


Figure 130–5—Transmit common-mode output return loss limit

130.7.1.7 Transition time

The rising and falling edge transition times shall be as specified in Table 130–4 when measured with respect to v_1 and v_3 as defined in Figure 130–7. Measurement is done using the square wave test pattern defined in 52.9.12, with no equalization (bit 1.150.3 in Table 45–69 set to ‘1’) and a run of at least eight consecutive ones followed by an equal number of consecutive zeros.

The BASE-R PMD control register is also used by 5GBASE-KR described in Clause 130 to disable the transmitter equalizer for test purposes. 5GBASE-KR does not use the start-up protocol.

130.7.1.8 Transmit jitter test requirements

Transmit jitter is defined with respect to a test procedure resulting in a BER bathtub curve such as that described in Annex 48B.3. For the purpose of jitter measurement, the effect of a single-pole high-pass filter with a 3 dB point at 4 MHz is applied to the jitter. The data pattern for jitter measurements shall be test

patterns 2 or 3 as defined in 52.9.1.1. Crossing times are defined with respect to the mid-point (0 V) of the AC-coupled differential signal. Equalization shall be off during jitter testing (bit 1.150.3 in Table 45–69 set to ‘1’).

The duty cycle distortion test pattern shall consist of alternating ones and zeros (i.e., 10101010...).

130.7.1.9 Transmit jitter

The transmitter shall be as specified in Table 130–4. Duty cycle distortion (DCD) is considered a component of deterministic jitters. The peak-to-peak duty cycle distortion is defined as the absolute value of the difference in the mean pulse width of a 1 pulse or the mean pulse width of a 0 pulse (as measured at the mean of the high- and low-voltage levels in a clock-like repeating 0101 bit sequence) and the nominal pulse width. Jitter specifications are specified for BER 10^{-12} . Transmit jitter test requirements are specified in 130.7.1.8.

NOTE—Duty cycle distortion (DCD) is also referred to as Even-odd jitter (see 92.8.3.8.1).

130.7.1.10 Transmitter output waveform

The 5GBASE-KR transmitter includes pre-emphasis to compensate for frequency-dependent loss in the backplane channel and facilitate data recovery at the receiver. This equalization may be accomplished with a two-tap finite impulse response (FIR) structure as shown in Figure 130–6.

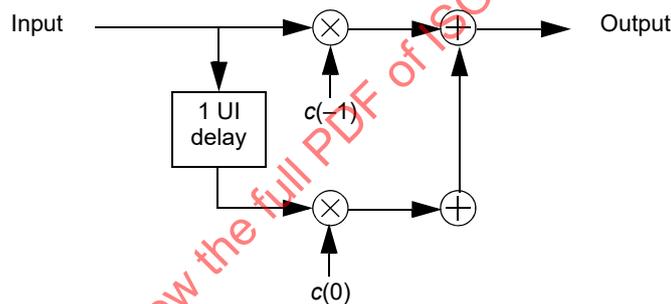


Figure 130–6—Transmit equalizer example

The test pattern for the transmitter output waveform is the square wave test pattern defined in 52.9.1.2, with a run of at least eight consecutive ones followed by an equal number of consecutive zeros. The transmitter output waveform test is based on the voltages v_1 through v_4 , which are illustrated in Figure 130–7 and described below. These measurements are used to calculate R_{pre} , defined in Equation (130–7). R_{pre} shall be within the limits specified in Table 130–4.

From these voltages, the pre-equalization ratio R_{pre} is derived from Equation (130–7).

$$R_{pre} = \frac{v_2}{v_1} \tag{130-7}$$

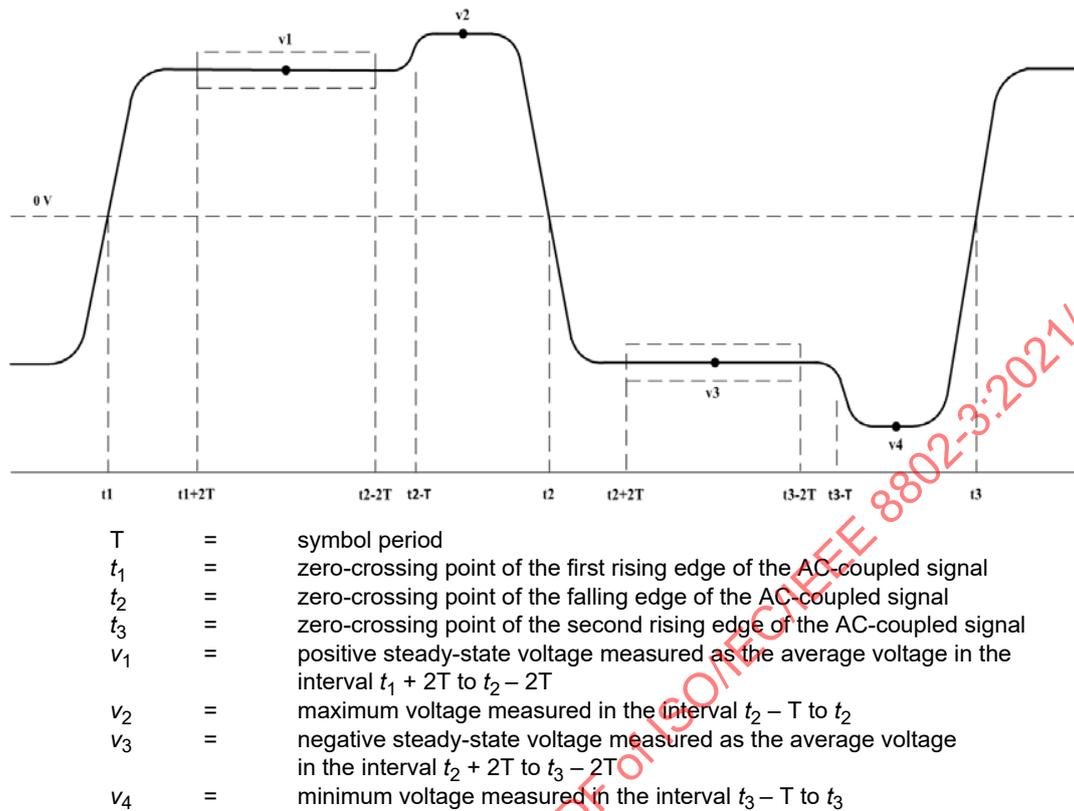


Figure 130-7—Transmitter output waveform

130.7.2 Receiver characteristics

Receiver characteristics at TP4 (see Figure 130-1) are summarized in Table 130-5 and detailed in 130.7.2.1 through 130.7.2.5.

Table 130-5—Receiver characteristics for 5GBASE-KR

Parameter	Subclause reference	Value	Units
Bit error ratio	130.7.2.1	10^{-12}	
Signaling speed	130.7.2.2	5.15625 ± 100 ppm	GBd
Differential input peak-to-peak amplitude (max)	130.7.2.4	1200 ^a	mV
Differential input return loss (min) ^b	130.7.2.5	Equation (130-3) and Equation (130-4)	dB

^aThe receiver shall tolerate input amplitudes up to a maximum of 1600 mV without permanent damage.

^bRelative to 100 Ω differential.

130.7.2.1 Receiver interference tolerance

The receiver interference tolerance consists of two separate tests as described in Annex 69A with the parameters specified in Table 130–6. The data pattern for the interference tolerance test shall be the test patterns 2 or 3 as defined in 52.9.1.1. The receiver shall satisfy the requirements for interference tolerance specified in Annex 69A for both tests.

Table 130–6—5GBASE-KR interference tolerance parameters

Parameter	Test 1 values	Test 2 values	Units
Target BER	10 ⁻¹²	10 ⁻¹²	
m _{TC} (min) ^a	1	0.5	
Amplitude of broadband noise (min RMS)	5.2	12	mV
Applied transition time (20%–80%, min)	60	60	ps
Applied Sinusoidal jitter (min peak-to-peak)	0.115	0.115	UI
Applied random jitter (min peak-to-peak) ^b	0.13	0.13	UI

^am_{TC} is defined in Equation 69A-5 in Annex 69A.
^bApplied random jitter is specified at a BER of 10⁻¹².

130.7.2.2 Signaling speed range

A signaling speed of a 5GBASE-KR receiver shall comply with the requirements of Table 130–5 for any signaling speed in the range 5.15625 GBd ± 100 ppm.

130.7.2.3 AC-coupling

The 5GBASE-KR receiver shall be AC-coupled to the backplane to allow for maximum interoperability between various 5 Gb/s components. AC-coupling is considered to be part of the receiver for the purposes of this specification unless explicitly stated otherwise. There may be various methods for AC-coupling in actual implementations.

NOTE—It is recommended that the maximum value of the coupling capacitors be limited to 100 nF. This will limit the inrush currents to the receiver that could damage the receiver circuits when repeatedly connected to transmit modules with a higher voltage level.

130.7.2.4 Input signal amplitude

5GBASE-KR receivers shall accept differential input signal peak-to-peak amplitudes produced by compliant transmitters connected without attenuation to the receiver, and still meet the BER requirement specified in 130.7.2.1. This may be larger than the differential maximum in 130.7.1.4 due to the actual transmitter output and receiver input impedances. The input impedance of a receiver can cause the minimum signal into a receiver to differ from that measured when the receiver is replaced with a 100 Ω test load. Since the channel is AC-coupled, the absolute voltage levels with respect to the receiver ground are dependent on the receiver implementation.

130.7.2.5 Differential input return loss

For frequencies from 100 MHz to 3750 MHz, the differential input return loss, in dB, of the receiver shall be as specified in Table 130–5.

130.8 Interconnect characteristics

Informative interconnect characteristics are provided in Annex 69B.

130.9 Environmental specifications

130.9.1 General safety

All equipment that meets the requirements of this standard shall conform to applicable sections (including isolation requirements) of IEC 60950-1.

130.9.2 Network safety

The designer is urged to consult the relevant local, national, and international safety regulations to ensure compliance with the appropriate requirements.

130.9.3 Installation and maintenance guidelines

It is recommended that sound installation practice, as defined by applicable local codes and regulations, be followed in every instance in which such practice is applicable.

130.9.4 Electromagnetic compatibility

A system integrating the 5GBASE-KR PHY shall comply with applicable local and national codes for the limitation of electromagnetic interference.

130.9.5 Temperature and humidity

A system integrating the 5GBASE-KR PHY is expected to operate over a reasonable range of environmental conditions related to temperature, humidity, and physical handling (such as shock and vibration). Specific requirements and values for these parameters are considered to be beyond the scope of this standard.

130.10 Protocol implementation conformance statement (PICS) proforma for Clause 130, Physical Medium Dependent (PMD) sublayer and baseband medium, type 5GBASE-KR⁸

130.10.1 Introduction

The supplier of a protocol implementation that is claimed to conform to IEEE Std 802.3, Clause 130, Physical Medium Dependent (PMD) sublayer and baseband medium type 5GBASE-KR, shall complete the following protocol implementation conformance statement (PICS) proforma. A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

130.10.2 Identification

130.10.2.1 Implementation identification

Supplier	
Contact point for inquiries about the PICS	
Implementation Name(s) and Version(s)	
Other information necessary for full identification—e.g., name(s) and version(s) for machines and/or operating systems; System Name(s)	
NOTE 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirements for the identification. NOTE 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier’s terminology (e.g., Type, Series, Model).	

130.10.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3-2018, Clause 130, Physical Medium Dependent (PMD) sublayer and baseband medium type 5GBASE-KR
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3-2018.)	
Date of Statement	

⁸Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

130.10.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
XGE	XGMII	130.1, Clause 46	Interface is supported	O	Yes [] No []
PCS	Support of 5GBASE-R PCS/PMA	130.1, Clause 49		M	Yes []
AN	Auto-Negotiation for Backplane Ethernet	130.1, Clause 73	Device implements Auto-Negotiation for Backplane Ethernet	M	Yes []
DC	Delay Constraints	130.4	Device conforms to delay constraints	M	Yes []
*LPI	LPI	130.6.10	LPI	O	Yes [] No []
*MD	MDIO interface	130.5	Device implements MDIO	O	Yes [] No []
*TD	Global_PMD_transmit_disable	130.6.5		O	Yes [] No []

**130.10.4 PICS proforma tables for Clause 130, Physical Medium Dependent (PMD) sublayer
and baseband medium, type 5GBASE-KR**

130.10.4.1 PCS requirements for AN service interface

Item	Feature	Subclause	Value/Comment	Status	Support
PR1	AN service interface primitive	130.3	The PCS associated with this PMD supports the primitive AN_LINK.indication defined in 73.9	M	Yes []

130.10.4.2 PMD functional specifications

Item	Feature	Subclause	Value/Comment	Status	Support
FS1	PMD combined with sublayers through interface defined in Clause 45.	130.1	5GBASE-KR PHY layer associations. See Table 130–1	M	Yes []
FS2	PMD_global_signal_detect set to the value of SIGNAL_DETECT	130.2.3.2	Required generation of PMD primitive	M	Yes []
FS3	Tx+Rx delays from 5GBASE-KR PMD and medium <= 1024 bit times	130.4		M	Yes []
FS4	PMD Transmit function	130.6.2	Conveys bits from PMD service interface to MDI	M	Yes []
FS5	PMD Transmitter signal	130.6.2	A positive differential voltage corresponds to tx_bit = ONE	M	Yes []
FS6	PMD Receive function	130.6.3	Conveys bits from MDI to PMD service interface	M	Yes []
FS7	PMD Receiver signal	130.6.3	A positive differential voltage corresponds to rx_bit = ONE	M	Yes []
FS8	PMD_fault indicates local fault in MDIO implementation	130.6.3	As specified in 45.2.1.7.4	MD:M	Yes [] N/A []
FS9	PMD Signal detect function	130.6.4	Report to PMD service interface	M	Yes []
FS10	Global signal detect	130.6.4	Value described in 45.2.1.9.7	M	Yes []
FS11	Global_PMD_signal_detect set to SIGNAL_DETECT	130.6.4	As described in Table 130–3	MD:M	Yes [] N/A []
FS12	SIGNAL_DETECT value	130.6.4	Set to FAIL	M	Yes []
FS13	PMD Signal detect during LPI	130.6.4	Detect signal energy during LPI	LPI:M	Yes [] N/A []
FS14	PMD Signal detect for EEE	130.6.4	Transition timing to set SIGNAL_DETECT	LPI:M	Yes [] N/A []
FS15	PMD Transmit disable requirements	130.6.5	Requirements of 130.6.5 and Table 130–4	TD:M	Yes [] N/A []
FS16	Loopback not affected by Global_PMD_transmit_disable	130.6.5		M	Yes []
FS17	PMD Transmit disable during LPI	130.6.5	Disable transmitter during tx_mode = QUIET	LPI:M	Yes [] N/A []
FS18	Loopback support	130.6.6	Provided for 5GBASE-KR PMD by transmitter and receiver	M	Yes []
FS19	Pre-cursor equalization ratio	130.7.1.10	As described in Table 130–4	M	Yes []

130.10.4.3 Management functions

Item	Feature	Subclause	Value/Comment	Status	Support
MF1	MDIO Variable Mapping	130.5	Per Table 130–2 and Table 130–3	MD:M	Yes [] N/A []
MF2	PMD_transmit_fault function	130.6.8	Sets PMD_transmit_fault as specified in 45.2.1.7.4	MD:M	Yes [] N/A []
MF3	PMD_receive_fault function	130.6.9	Sets PMD_receive_fault as specified in 45.2.1.7.5	MD:M	Yes [] N/A []

130.10.4.4 PMD Transmitter electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
TC1	Test fixture impedance	130.7.1.2	100 Ω	M	Yes []
TC2	Differential output return loss of test fixture	130.7.1.2	Per Equation (130–1) and Equation (130–2)	M	Yes []
TC3	Signaling speed	130.7.1.3	5.15625 GBd ± 100 ppm	M	Yes []
TC4	Transmitter differential peak-to-peak voltage	130.7.1.4	1200 mV, pk-pk, using the pattern as defined in 130.7.1.4	M	Yes []
TC5	Maximum transmitter differential peak-to-peak voltage when TX disabled	130.7.1.4	30 mV	M	Yes []
TC6	Common-mode output voltage	130.7.1.4	0V to +1.9V	M	Yes []
TC7	Tx output to DATA compliance timing	130.7.1.4	Fully compliant ≤ 5 μsec from tx_mode = DATA	M	Yes []
TC8	Output Amplitude LPI voltage	130.7.1.4	Less than or equal to 30 mV within 500 ns of tx_quiet	LPI:M	Yes [] N/A []
TC9	Output Amplitude ON voltage	130.7.1.4	Less than or equal to 720 mV within 500 ns of tx_quiet deasserted	LPI:M	Yes [] N/A []
TC10	Common-mode maintained within voltage limits	130.7.1.4	Output within ± 150 mV of the pre-LPI value	LPI:M	Yes [] N/A []
TC11	Differential output return loss	130.7.1.5	Per Equation (130–3) and Equation (130–4)	M	Yes []
TC12	Differential output reference impedance	130.7.1.5	100 Ω	M	Yes []
TC13	Common-mode output return loss	130.7.1.6	Per Equation (130–5) and Equation (130–6)	M	Yes []
TC14	Rising edge transition time	130.7.1.7	Between 20 ps and 60 ps measured at the 20% and 80% of peak-to-peak differential with pattern in 52.9.1.2	M	Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
TC15	Falling edge transition time	130.7.1.7	Between 20 ps and 60 ps measured at the 20% and 80% of peak-to-peak differential with pattern in 52.9.1.2	M	Yes []
TC16	Transmit jitter, peak-to-peak	130.7.1.9	Max Tj of 0.27 UI	M	Yes []
TC17	Duty Cycle Distortion	130.7.1.9	Not to exceed 0.035 UI	M	Yes []
TC18	Jitter test patterns	130.7.1.8	Test patterns 2 or 3 as defined in 52.9.1.1	M	Yes []
TC19	DCD test pattern	130.7.1.8	Test patterns 2 or 3 as defined in 130.7.1.8.	M	Yes []
TC20	Tx output waveform measurement criteria: V ₁ , V ₂ , V ₃ , V ₄	130.7.1.10	Measured as per Figure 130–7.	M	Yes []

130.10.4.5 Receiver electrical characteristics

Item	Feature	Subclause	Value/Comment	Status	Support
RC1	Receiver amplitude tolerance	130.7.2	Amplitudes of ≤1600 mV without permanent damage	M	Yes []
RC2	Receiver interference tolerance	130.7.2.1	Measured as described in Annex 69A with parameters in Table 130–6	M	Yes []
RC3	Receiver interference tolerance	130.7.2.1	Test patterns 2 or 3 as defined in 59.9.1.1	M	Yes []
RC4	Receiver interference tolerance	130.7.2.1	Satisfy the requirements specified in Annex 69A	M	Yes []
RC5	Signaling speed	130.7.2.2	5.15625 GBd ±100 ppm	M	Yes []
RC6	Receiver coupling	130.7.2.3	AC-coupled	M	Yes []
RC7	Input signal amplitude	130.7.2.4	BER met when compliant transmitter is connected with no attenuation	M	Yes []
RC8	Differential input return loss	130.7.2.5	Per Equation (130–3) and Equation (130–4)	M	Yes []
RC9	Differential input return loss reference impedance	130.7.2.5	100 Ω	M	Yes []

130.10.4.6 Environmental specifications

Item	Feature	Subclause	Value/Comment	Status	Support
ES1	General safety	130.9.1	Complies with applicable section of IEC 60950-1	M	Yes []
ES2	Electromagnetic interference	130.9.4	Complies with applicable local and national codes	M	Yes []

Annex 31B

(normative)

MAC Control PAUSE operation

31B.3 Detailed specification of PAUSE operation

31B.3.7 Timing considerations for PAUSE operation

Change fifth and sixth paragraphs in 31B.3.7 as shown:

At operating speeds of 2.5 Gb/s, a station with a 2.5GBASE-T PHY shall not begin to transmit a (new) frame more than 34 pause_quanta after the reception of a valid PAUSE frame that contains a non-zero value of pause_time, as measured at the MDI. A station using any other 2.5 Gb/s PHY shall not begin to transmit a (new) frame more than 6 pause_quanta after the reception of a valid PAUSE frame that contains a non-zero value of pause_time, as measured at the MDI.

At operating speeds of 5 Gb/s, a station with a 5GBASE-T PHY shall not begin to transmit a (new) frame more than 42 pause_quanta after the reception of a valid PAUSE frame that contains a non-zero value of pause_time, as measured at the MDI. A station using any other 5 Gb/s PHY shall not begin to transmit a (new) frame more than 12 pause_quanta after the reception of a valid PAUSE frame that contains a non-zero value of pause_time, as measured at the MDI.

Change the beginning of the list after the last paragraph (“The PAUSE response time”) in 31B.3.7 as shown:

2.5 Gb/s (using 2.5GBASE-T) – max_overrun = 2176 + frame_length

2.5 Gb/s (not using 2.5GBASE-T) – max_overrun = 384 + frame_length

5 Gb/s (using 5GBASE-T) – max_overrun = 2624 + frame_length

5 Gb/s (not using 5GBASE-T) – max_overrun = 768 + frame_length

31B.4 Protocol implementation conformance statement (PICS) proforma for MAC Control PAUSE operation⁹

31B.4.3 Major capabilities/options

Change the table in 31B.4.3 by adding the following new rows for *MIIda and *MIIEa as shown (other unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
*MIIda	At operating speeds of 2.5 Gb/s with PHY types of 2.5GBASE-T.	31B.3.7	N/A	<u>O</u>	Yes [] No []
*MIId	At operating speeds of 2.5 Gb/s with PHY types other than 2.5GBASE-T.	31B.3.7	N/A	O	Yes [] No []
*MIIEa	At operating speeds of 5 Gb/s with PHY types other than 5GBASE-T.	31B.3.7	N/A	<u>O</u>	Yes [] No []
*MIIE	At operating speeds of 5 Gb/s with PHY types of 5GBASE-T.	31B.3.7	N/A	O	Yes [] No []

31B.4.6 PAUSE command MAC timing considerations

Change the table in 31B.4.6 by adding the following new rows for TIM4a and TIM5a as shown (other unchanged rows not shown):

Item	Feature	Subclause	Value/Comment	Status	Support
TIM4a	Measurement point for station at 2.5 Gb/s with PHY type other than 2.5GBASE-T	31B.3.7	Delay at MDI \leq 6 pause_quanta ^a	MIIEa: M	Yes [] N/A []
TIM5	Measurement point for station at 2.5 Gb/s with PHY type of 2.5GBASE-T	31B.3.7	Delay at MDI \leq 34 pause_quanta ^a	MIId: M	Yes [] N/A []
TIM5a	Measurement point for station at 5 Gb/s with PHY type other than 5GBASE-T	31B.3.7	Delay at MDI \leq 12 pause_quanta ^a	MIIEa: M	Yes [] N/A []
TIM6	Measurement point for station at 5 Gb/s with PHY type of 5GBASE-T	31B.3.7	Delay at MDI \leq 42 pause_quanta ^a	MIIE: M	Yes [] N/A []

^aDelay from receiving valid PAUSE command, with nonzero value for pause_time, to cessation of transmission.

⁹Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

Annex 69A

(normative)

Interference tolerance testing

69A.1 Introduction

Change the last paragraph of 69A.1 as shown:

For the channel to work, the receiver must be able to extract correct data from the lossy channel in the ~~presences~~ presence of interference. The ability of the receiver to extract data in the presence of interference is an important characteristic of the receiver and needs to be measured. This ability is called interference tolerance.

69A.2 Test setup

69A.2.1 Pattern generator

Change the beginning of 69A.2.1 as shown:

For 1000BASE-KX and 10GBASE-KX4, the amplitude delivered by the pattern generator to the test channel shall be no greater than the specified minimum transmitter output amplitude for the port type being tested adjusted by a gain b_{TC} as defined in 69A.2.2.

For 2.5GBASE-KX, the peak-to-peak amplitude delivered by the pattern generator, as measured on a sequence of alternating ones and zeros, shall be no more than 800 mV, adjusted by a gain b_{TC} as defined in 69A.2.2.

For 5GBASE-KR, 10GBASE-KR, and 40GBASE-KR4, the peak-to-peak amplitude delivered by the pattern generator, as measured on a sequence of alternating ones and zeros, shall be no more than 800 mV, adjusted by a gain b_{TC} as defined in 69A.2.2, regardless of equalization setting.

Change the last paragraph in 69A.2.1 as shown:

The pattern generator may include equalization depending on the port type being tested. For 1000BASE-KX and 2.5GBASE-KX, the pattern generator shall not include equalization. For 5GBASE-KR, equalization equivalent to a two-tap transversal filter meeting the requirements of Table 130-4 shall be included. For 10GBASE-KX4, the pattern generator shall include equalization such that the differential output template defined in 71.7.1.6 is met. For 10GBASE-KR, equalization equivalent to a three-tap transversal filter meeting the requirements of 72.7.1.10 shall be included.

69A.2.2 Test channel

Change the last two last paragraphs in 69A.2.2 as shown:

The values f_1 and f_2 are a function of the port type under test (see Table 69B-1 and Table 69B-2), and A_{\max} is defined in 69B.4.2.

The test channel shall have m_{TC} greater than the minimum value specified for the port type under test and the test being performed. The test channel return loss, as measured at TP1 (see Figure 69A-1) and TP4 (see Figure 69A-1), shall be greater than or equal to 20 dB from f_{\min} to f_2 .

69A.2.4 Transmitter control

Change 69A.2.4 as shown:

For 10GBASE-KR testing, the pattern generator is controlled by transmitter control. Transmitter control responds to inputs from the receiver to adjust the equalization of the pattern generator. The receiver may communicate through its associated transmitter, using the protocol described in 72.6.10, or by other means. For 2.5GBASE-KX and 5GBASE-KR testing, transmitter control is not used.

69A.3 Test methodology

Change 69A.3 as shown:

For 10GBASE-KR testing, the pattern generator shall first be configured to transmit the training pattern defined in 72.6.10.2. During this initialization period, the DUT shall configure the pattern generator equalizer, via transmitter control, to the coefficient settings it would select using the protocol described in 72.6.10. During training, the broadband noise measured at TP4 (see Figure 69A-1) shall have RMS amplitude less than 1 mV. Training patterns and transmitter control are not used for 2.5GBASE-KX and 5GBASE-KR testing.

The pattern generator shall be configured to transmit the test pattern defined for the port type under test.

The broadband noise source shall then be set to the amplitude specified for the port type being tested, as measured at TP4 (see Figure 69A-1). The measured BER shall be less than the target BER specified for the port type under test.

The interference tolerance test parameters are specified in Table 70-7 for 1000BASE-KX, in Table 71-7 for 10GBASE-KX4, in Table 128-6 for 2.5GBASE-KX, in Table 130-6 for 5GBASE-KR, and in Table 72-10 for 10GBASE-KR and 40GBASE-KR4.

Annex 69B

(informative)

Interconnect characteristics

69B.4 Channel parameters

69B.4.1 Overview

Change the third paragraph of 69B.4.1 as shown:

The informative parameters for channel insertion loss are summarized in Table 69B-1 and Table 69B-2.

Change the title of Table 69B-1 as shown:

Table 69B-1—Insertion loss parameters for 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR4

Insert the following new table (Table 69B-2) after Table 69B-1:

Table 69B-2—Insertion loss parameters for 2.5GBASE-KX and 5GBASE-KR

Parameter	2.5GBASE-KX	5GBASE-KR	Units
f_{min}	0.05		GHz
f_{max}	2.34375	3.8671875	GHz
b_1	2.4×10^{-5}		
b_2	1.3×10^{-10}		
b_3	4.9×10^{-20}		
b_4	-1.9×10^{-30}		
f_1	0.312	0.5	GHz
f_2	1.5625	2.578125	GHz
f_a	0.1		GHz
f_b	1.5625	2.578125	GHz

69B.4.2 Fitted attenuation

Change the last paragraph of 69B.4.2 as shown:

It is recommended that the fitted attenuation of the channel be less than or equal to A_{max} as defined by Equation (69B-6), where f is expressed in Hz and the coefficients b_1 through b_4 are given in Table 69B-1 and Table 69B-2.

$$A(f) \leq A_{max}(f) = 20 \log_{10}(e) \times (b_1 \sqrt{f} + b_2 f + b_3 f^2 + b_4 f^3) \quad (69B-6)$$

for $f_1 \leq f \leq f_2$. The fitted attenuation limit is illustrated in Figure 69B-2 and Figure 69B-2a.

Change the title of Figure 69B–2 as shown (the figure art remains unchanged):

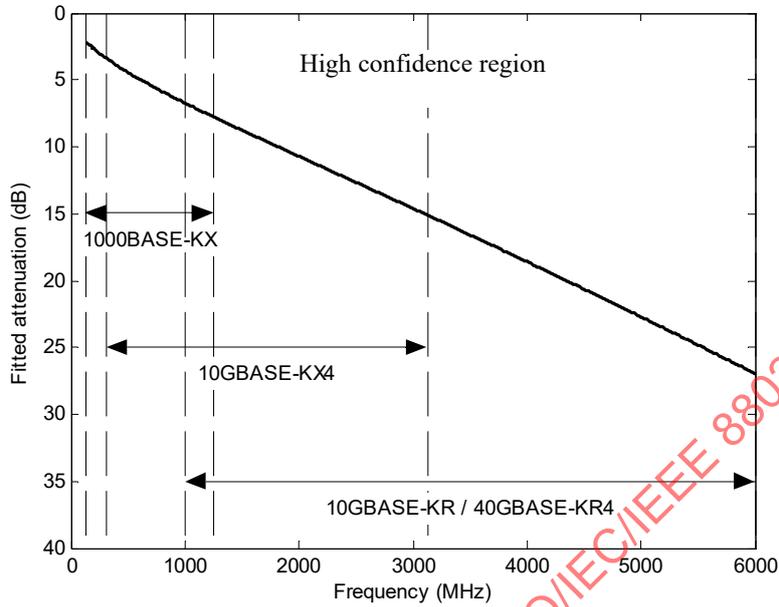


Figure 69B–2—Fitted attenuation limit for 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR4

Insert the following new figure (Figure 69B–2a) after Figure 69B–2:

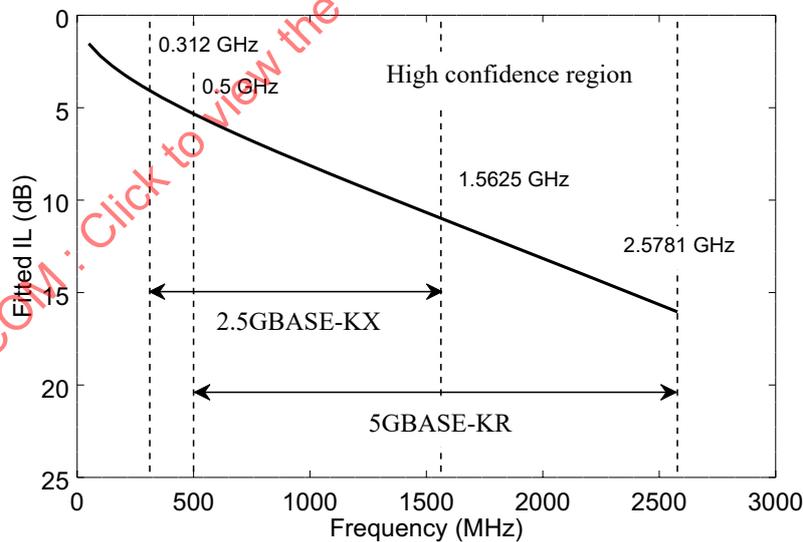


Figure 69B–2a—Fitted attenuation limit for 2.5GBASE-KX and 5GBASE-KR

69B.4.3 Insertion loss

Change the text of 69B.4.3 as shown (Figure 69B-3, Figure 69B-4, and Figure 69B-5 remain unchanged):

Insertion loss is defined as the magnitude, expressed in decibels, of the differential response measured from TP1 to TP4. For 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR, it is recommended that the insertion loss magnitude, IL , be within the high confidence region defined by Equation (69B-7) and Equation (69B-8).

$$IL(f) \leq IL_{\max}(f) = A_{\max}(f) + 0.8 + 2.0 \times 10^{-10} f \quad (69B-7)$$

for $f_{\min} \leq f \leq f_2$

$$IL(f) \leq IL_{\max}(f) = A_{\max}(f) + 0.8 + 2.0 \times 10^{-10} f_2 + 1 \times 10^{-8} (f - f_2) \quad (69B-8)$$

for $f_2 < f \leq f_{\max}$

For 2.5GBASE-KX, it is recommended that the insertion loss magnitude, IL , be within the high confidence region defined by Equation (69B-8a).

$$IL(f) \leq \left\{ \begin{array}{ll} 0.668 + 3.755\sqrt{f} + 3.608f & f_{\min} \leq f < f_2 \\ -23.753 + 22.242f & f_2 \leq f < f_{\max} \end{array} \right\} \text{ (dB)} \quad (69B-8a)$$

For 5GBASE-KR, it is recommended that the insertion loss magnitude, IL , be within the high confidence region defined by Equation (69B-8b).

$$IL(f) \leq \left\{ \begin{array}{ll} 0.668 + 3.755\sqrt{f} + 3.608f & f_{\min} \leq f < f_2 \\ -18.753 + 13.48f & f_2 \leq f < f_{\max} \end{array} \right\} \text{ (dB)} \quad (69B-8b)$$

The values of f_{\min} , f_2 , and f_{\max} are given in Table 69B-1 and Table 69B-2 and A_{\max} is given in Equation (69B-6). The insertion loss limit is illustrated in Figure 69B-3, Figure 69B-4 and Figure 69B-5 through Figure 69B-5b.

Insert the following new figures (Figure 69B-5a and Figure 69B-5b) after Figure 69B-5:

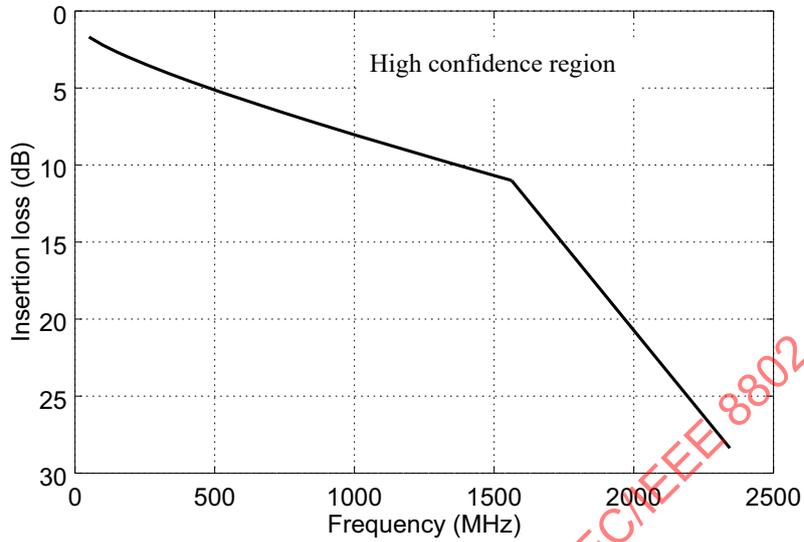


Figure 69B-5a—Insertion loss limit for 2.5GBASE-KX

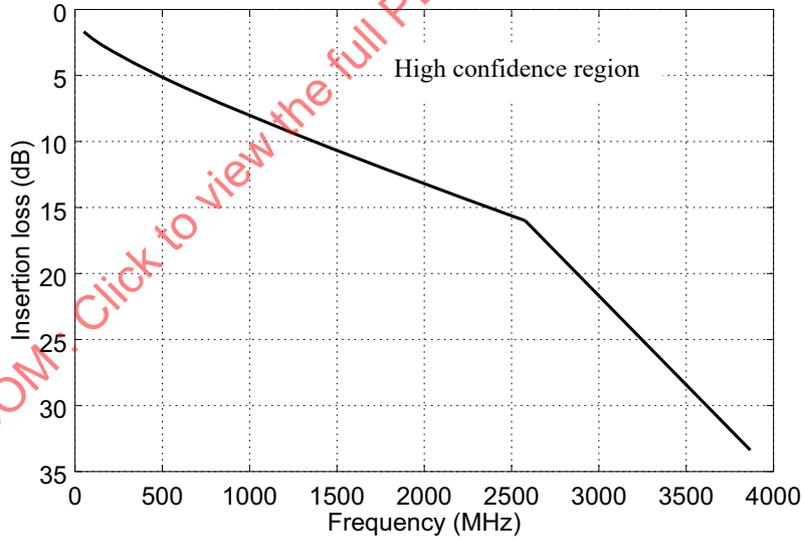


Figure 69B-5b—Insertion loss limit for 5GBASE-KR

69B.4.4 Insertion loss deviation

Change the text of 69B.4.4 as shown:

The insertion loss deviation, as defined by Equation (69B-9), is the difference between the insertion loss and the fitted attenuation defined in 69B.4.2.

$$ILD(f) = IL(f) - A(f) \quad (69B-9)$$

For 100BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR, it is recommended that ILD be within the high confidence region defined by the following equations: Equation (69B-10) and Equation (69B-11).

$$ILD(f) \geq ILD_{min}(f) = -1.0 - 0.5 \times 10^{-9} f \quad (69B-10)$$

$$ILD(f) \leq ILD_{max}(f) = 1.0 + 0.5 \times 10^{-9} f \quad (69B-11)$$

for $f_1 \leq f \leq f_2$

For 2.5GBASE-KX and 5GBASE-KR, it is recommended that ILD be within the high confidence region defined by Equation (69B-11a) and Equation (69B-11b).

$$ILD(f) \geq ILD_{min}(f) = -1.0 - 0.7 \times 10^{-9} f \quad (69B-11a)$$

$$ILD(f) \leq ILD_{max}(f) = 1.0 + 0.7 \times 10^{-9} f \quad (69B-11b)$$

for $f_1 \leq f \leq f_2$.

The recommendation applies over the frequency range $f_1 \leq f \leq f_2$. The values of f_1 and f_2 are dependent on port type and are given in Table 69B-1 and Table 69B-2. The insertion loss deviation limits for each port type is illustrated in Figure 69B-6 and Figure 69B-6a.

Change the title of Figure 69B-6 as shown (the figure art remains unchanged):

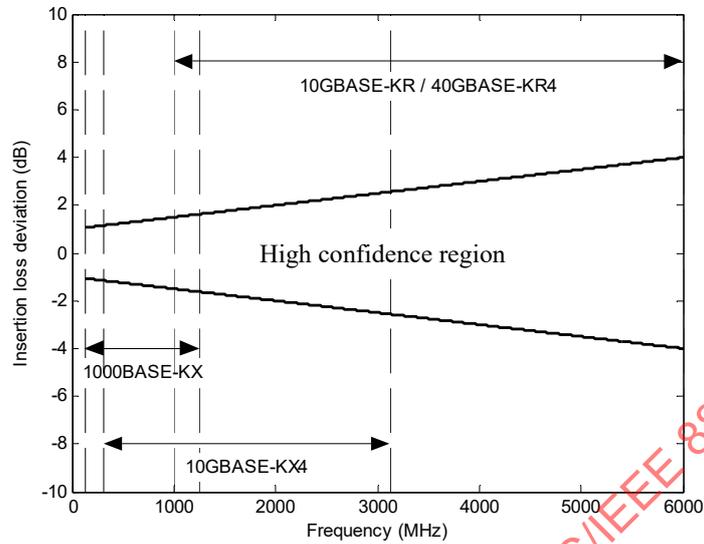


Figure 69B-6—Insertion loss deviation limits for 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR4

Insert the following new figure (Figure 69B-6a) after Figure 69B-6:

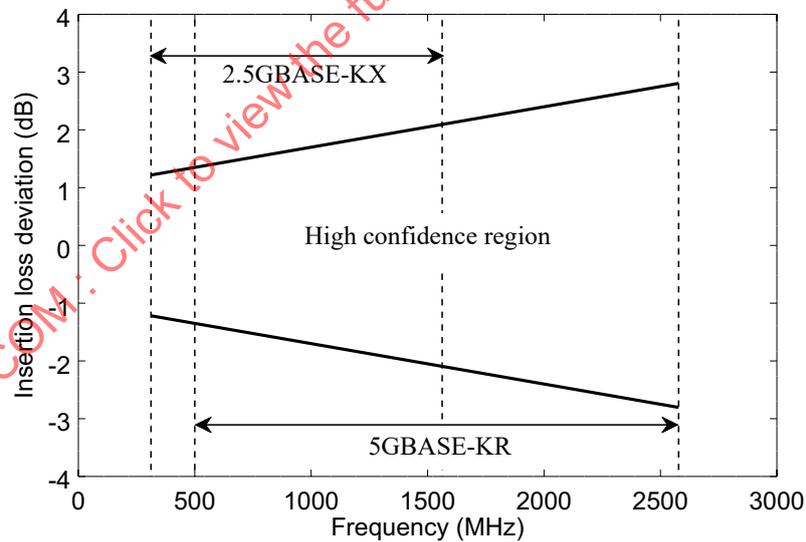


Figure 69B-6a—Insertion loss deviation limits for 2.5GBASE-KX and 5GBASE-KR

69B.4.5 Return loss

Change the last paragraph of 69B.4.5 as shown:

The recommendation applies from 50 MHz to the signaling speed of the PHY type of interest. The return loss limit is illustrated in Figure 69B-7 and Figure 69B-7a.

Change the title of Figure 69B-7 as shown (the figure art remains unchanged):

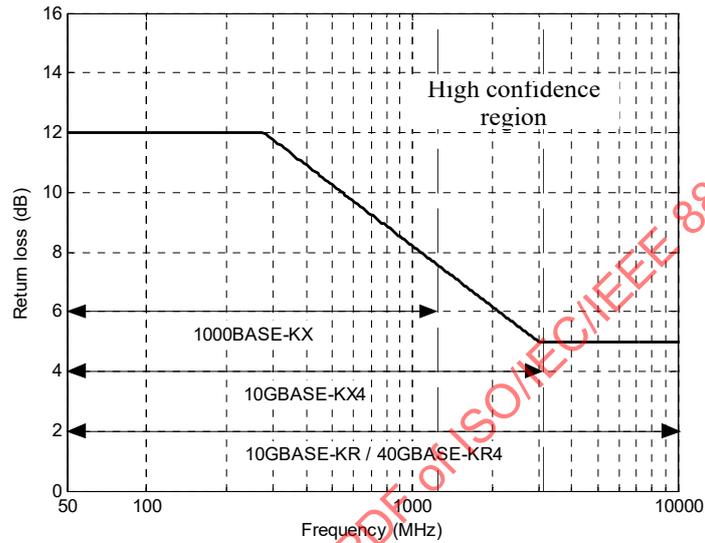


Figure 69B-7—Return loss limit for 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR4

Insert the following new figure (Figure 69B-7a) after Figure 69B-7:

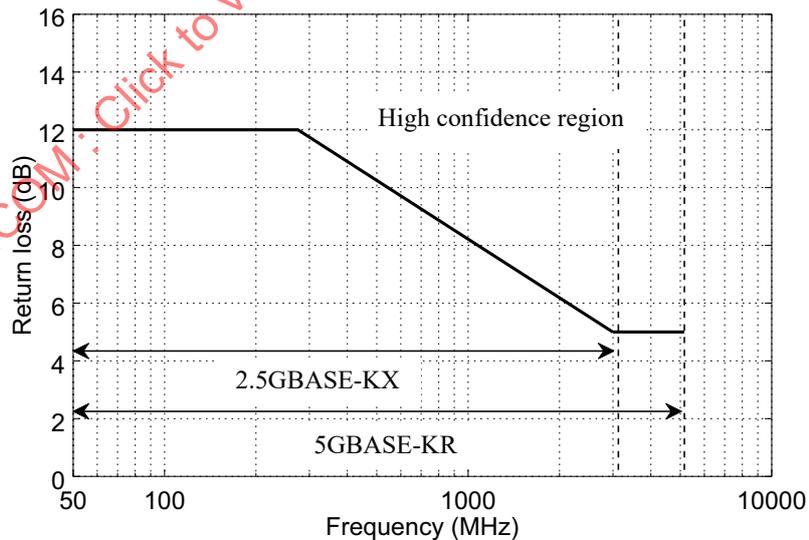


Figure 69B-7a—Return loss limit for 2.5GBASE-KX and 5GBASE-KR

69B.4.6 Crosstalk

69B.4.6.4 Insertion loss to crosstalk ratio (ICR)

Change the final paragraphs of 69B.4.6.4 as shown:

It is recommended that ICR_{fit} be greater than or equal to ICR_{min} as defined by the following equation: For 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR, ICR_{min} is defined by Equation (69B-24). For 2.5GBASE-KX and 5GBASE-KR, ICR_{min} is defined by Equation (69B-25).

$$ICR_{fit}(f) \geq ICR_{min}(f) = 23.3 - 18.7 \log_{10} \left(\frac{f}{5 \text{ GHz}} \right) \tag{69B-24}$$

$$ICR_{fit}(f) \geq ICR_{min}(f) = 28 - 18 \log_{10} \left(\frac{f}{2 \text{ GHz}} \right) \tag{69B-25}$$

ICR_{min} is defined over the frequency range for $f_a \leq f \leq f_b$. ICR_{fit} accounts for the worst-case differences in characteristics (e.g., amplitude, transition times) between the victim and aggressor transmitters. It also assumes a 3 dB signal-to-noise ratio penalty related to insertion loss deviation.

The insertion loss to crosstalk ratio limit for each port type is illustrated in Figure 69B-8 and Figure 69B-9.

Change the title of Figure 69B-8 as shown (the figure art remains unchanged):

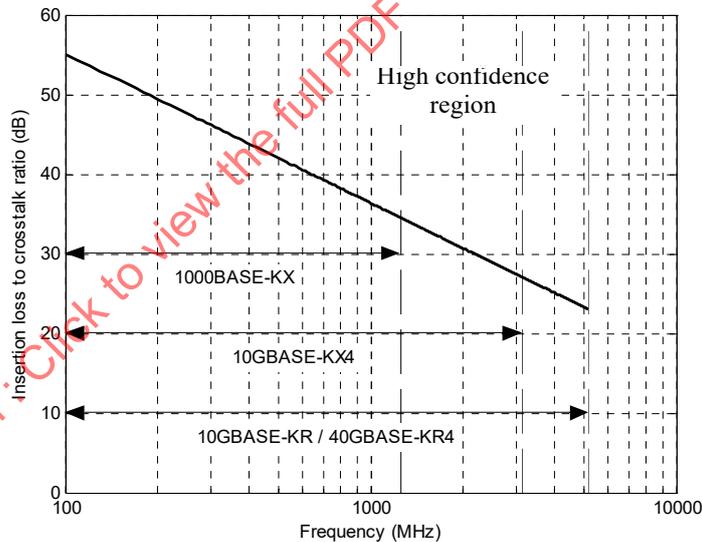


Figure 69B-8—Insertion loss to crosstalk ratio limit for 1000BASE-KX, 10GBASE-KX4, 10GBASE-KR, and 40GBASE-KR4

Insert the following new figure (Figure 69B-9) after Figure 69B-8:

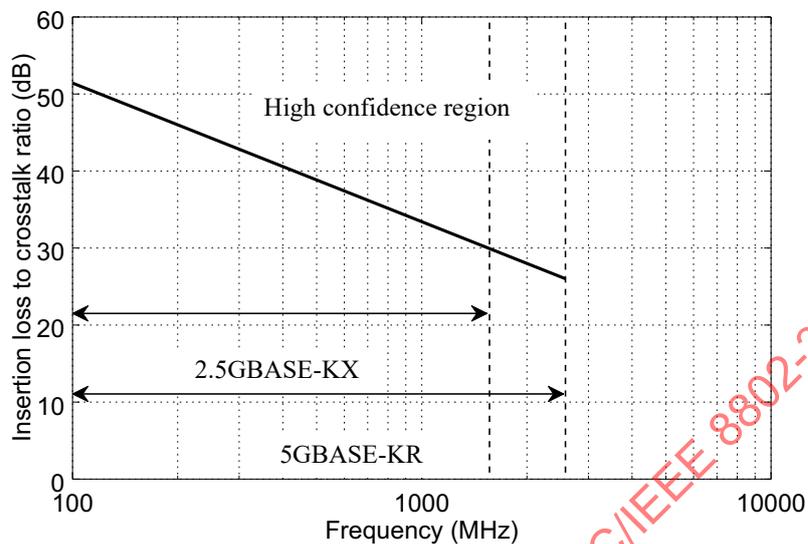


Figure 69B-9—Insertion loss to crosstalk ratio limit for 2.5GBASE-KX and 5GBASE-KR

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

Insert the following new annexes (Annex 127A, Annex 128A, Annex 128B, and Annex 130A) in alphanumeric order (see earlier in this amendment for the addition of corresponding clauses):

Annex 127A

(informative)

Compatibility of 2.5GBASE-X PCS/PMA with 1000BASE-X PCS/PMA running 2.5 times faster

This annex discusses the restrictions when operating 2.5GBASE-X PCS/PMA with a 1000BASE-X PCS/PMA link partner running 2.5 times faster. Compatibility of the PMD is outside the scope of this annex. In this annex when 1000BASE-X PCS/PMA is referred to, the 2.5 times speed up is implied.

The 2.5GBASE-X PCS/PMA is specified to be compatible with a link partner running 1000BASE-X PCS/PMA running 2.5 times faster as shown in Figure 127A–1 with some restrictions.

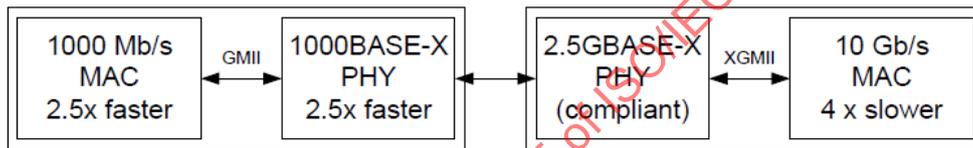


Figure 127A–1—2.5X speed 1000BASE-X to 2.5GBASE-X Link

2.5GBASE-X is defined to operate only in full duplex. Hence the 1000BASE-X PCS can only operate in full duplex. If the half duplex carrier extend or carrier extend with errors is sent by the 1000BASE-X PCS, the 2.5GBASE-X PCS will convert it to receive errors. A /T/R/ or a /T/R/R/ at the end of packet are correctly converted as idles.

The 2.5GBASE-X PCS does not support Clause 37 Auto-Negotiation. Hence, the 1000BASE-X PCS is expected to have its Clause 37 Auto-Negotiation functionality disabled so that the /C/ ordered set will not be transmitted. If a 2.5GBASE-X PCS receives /C/ ordered set, then undefined behavior may occur.

Since the 2.5GBASE-X PCS is attached to an RS that can send out sequence ordered set (/Q/), a compliant 1000BASE-X PCS will interpret each /Q/ ordered set as four /I/ ordered sets. Unlike /I/, /LI/, and /C/ ordered sets, there is no concept of correcting vs preserving the running disparity when /Q/ ordered sets are generated.

Unlike the GMII, there is no false carrier defined in the XGMII. The 2.5GBASE-X PCS Receive process can detect false carriers, but these will be converted to receive errors by the PCS Word Decode process.

It is permissible for a compliant 1000BASE-X PCS transmit process to truncate the first byte of a preamble in order to align the start of packet on the EVEN boundary. The implication of this is the 2.5 Gb/s RS (see 46.3.3.3) has to be able to accept a seven byte preamble on the XGMII with the SFD positioned on lane 2.

Since the start of packet has to align to lane 0 of the XGMII the 2.5 Gb/s RS may delete idle bytes in order to do the alignment. The 1G RS running 2.5 times faster has to be able to handle IPG shrinkage.

Annex 128A

(normative)

2.5 Gb/s Storage Enclosure Interface (2.5GSEI)

128A.1 Overview

This clause defines the functional and electrical characteristics for 2.5GSEI. This interface provides electrical characteristics and associated compliance points, which can optionally be used when designing systems with pluggable storage drive module interfaces. The compliance point definitions provide a unique partitioning of the channel defined in Annex 69B, such that the test points TP0_{D-H} and TP0_{H-D} defined in this annex are equivalent to TP1, and TP5_{D-H} and TP5_{H-D} are equivalent to TP4. Figure 128A-1 shows the test point locations associated with 2.5GSEI.

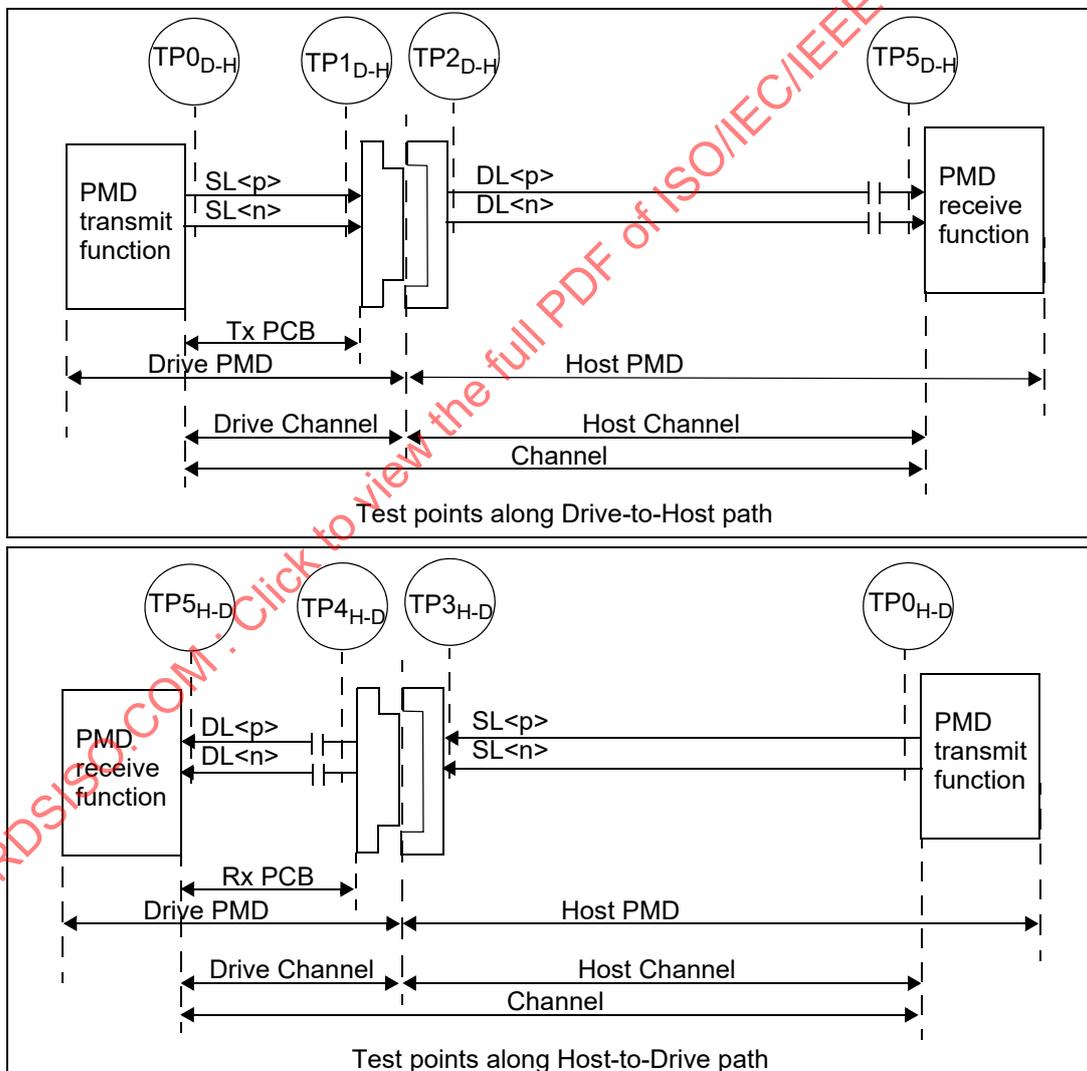
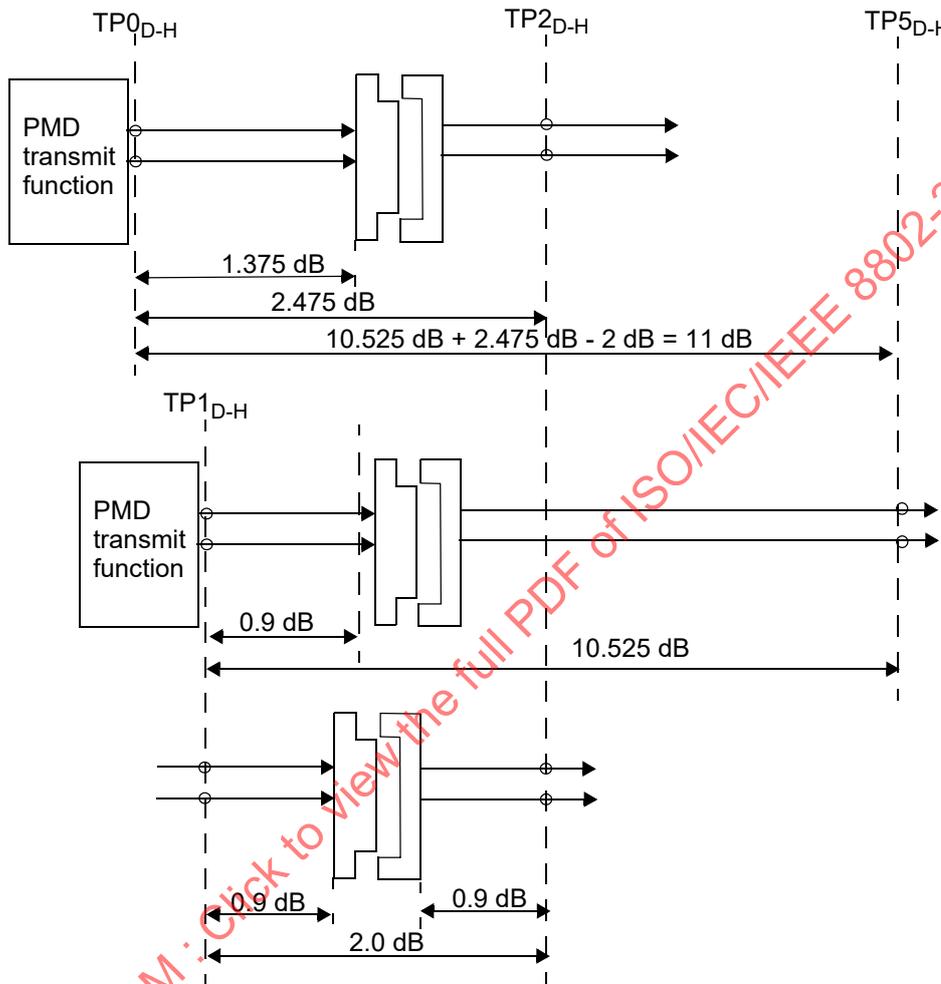


Figure 128A-1—Test points

The 2.5GSEI link is described in terms of a host 2.5GSEI component with associated insertion loss and a drive 2.5GSEI component. Figure 128A–2 (one direction shown) depicts the informative differential insertion loss budget at 1.5625 GHz for a typical 2.5GSEI application. The informative maximum differential insertion loss from TP0 to TP5 is given in Equation (128A–1) and depicted in Figure 128A–3. The 2.5GSEI interface consists of independent data paths in each direction. Each data path contains one differential lane, which is AC-coupled on the receiver side. The nominal signaling rate for each lane is 3.125 GBd.



NOTE—The connector insertion loss is 0.2 dB for the mated test fixture.

Figure 128A–2—Insertion loss budget at 1.5625 GHz

$$Insertion_loss(f) \leq \left\{ \begin{array}{ll} 0.668 + 3.755\sqrt{f} + 3.608f & 0.05 \leq f < 1.5625 \\ -23.753 + 22.242f & 1.5625 \leq f < 2.34375 \end{array} \right\} \text{ (dB)} \quad (128A-1)$$

where

f is the frequency in GHz

$Insertion_loss(f)$ is chip to chip (C2C) insertion loss

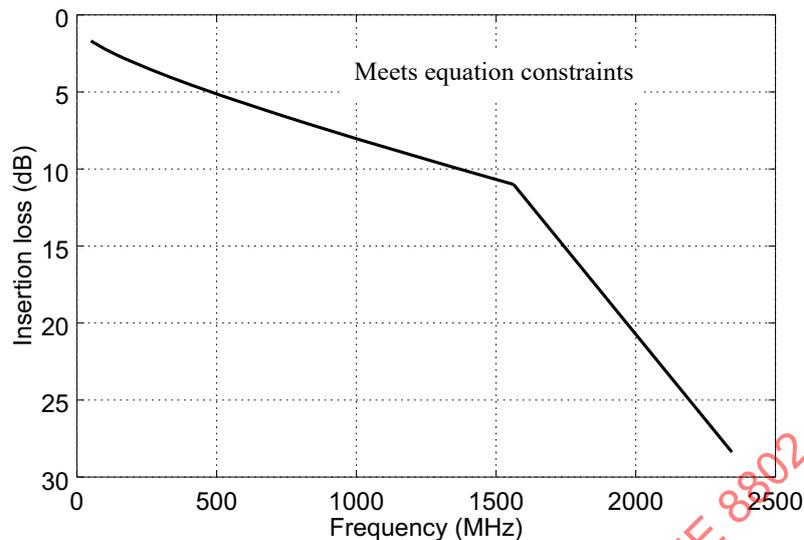


Figure 128A-3—Informative maximum differential insertion loss from TP0 to TP5

128A.1.1 Bit error ratio

The bit error ratio (BER) shall be less than 10^{-12} with any errors sufficiently uncorrelated to ensure an acceptably high mean time to false packet acceptance (MTTFA) assuming 8B/10B coding.

128A.2 2.5GSEI compliance point definitions

The electrical characteristics for 2.5GSEI are defined at compliance points for the host and drive, respectively. Reference test fixtures, called compliance boards, are used to access the electrical specification parameters. Figure 128A-4 depicts the location of compliance points when measuring host 2.5GSEI compliance. The output of the Host Compliance Board (HCB) is used to verify the host electrical output signal at TP4_{H-D} (see Figure 128A-1). Similarly, the input of the HCB at TP1_{D-H} (see Figure 128A-2) is used to verify the host input compliance.

Figure 128A-5 depicts the location of compliance points when measuring drive 2.5GSEI compliance. The output of the Drive Compliance Board (DCB) is used to verify the Drive electrical output signal at TP2_{D-H} (see Figure 128A-5). Similarly, the input of the DCB at TP3_{H-D} (see Figure 128A-5) is used to verify the Drive input compliance. Additional details on the requirements for the HCB and DCB are given in Annex 128B.

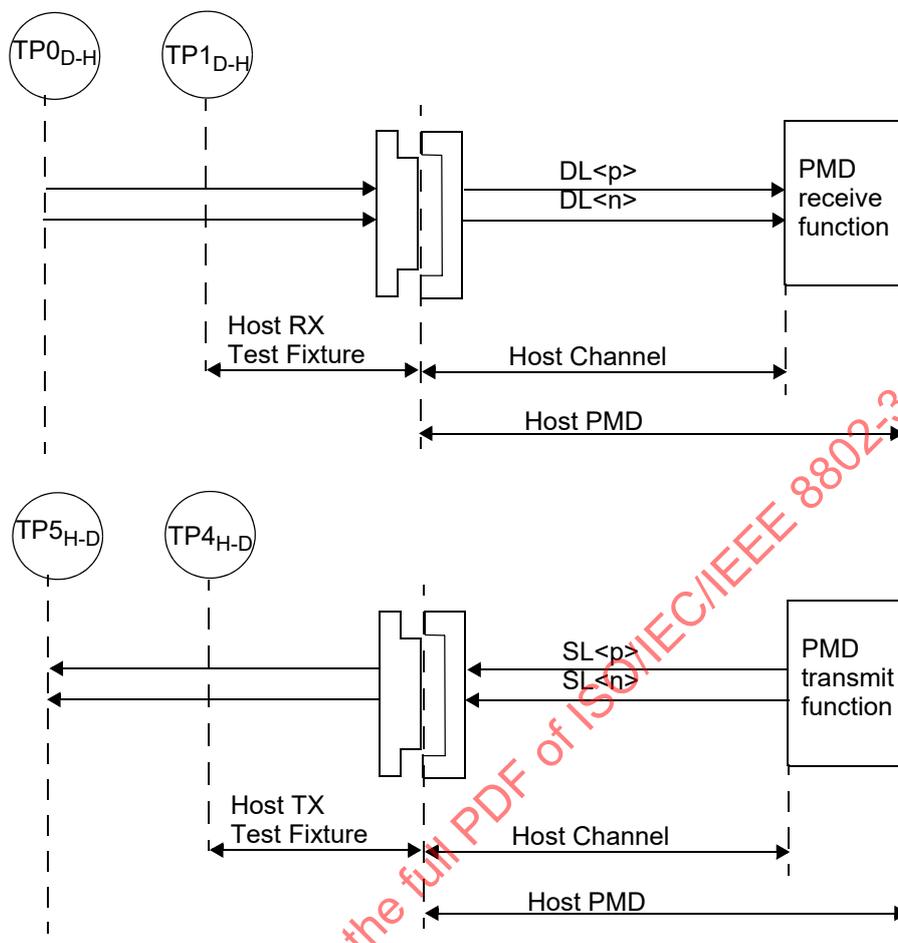


Figure 128A-4—Host compliance board

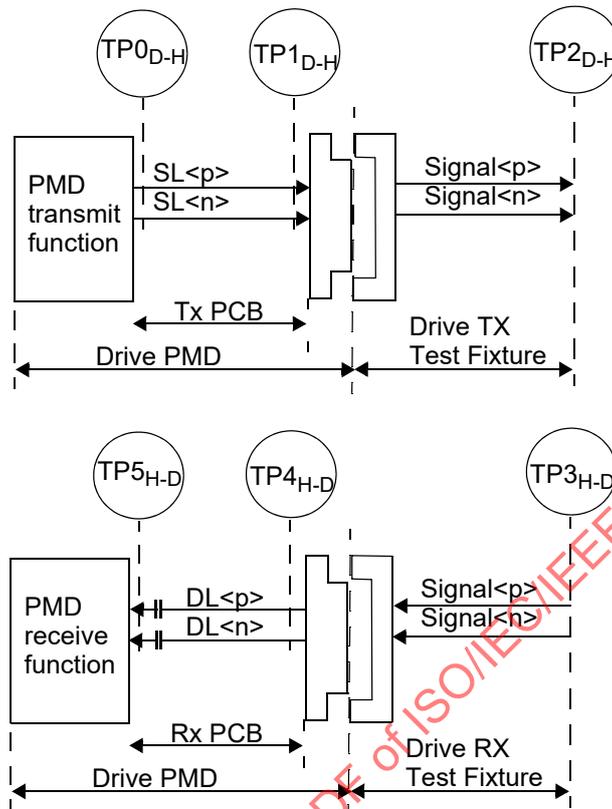


Figure 128A-5—Drive compliance board

STANDARDSISO.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 1:2021

128A.3 2.5GSEI electrical characteristics

128A.3.1 2.5GSEI host output characteristics

A 2.5GSEI host output shall meet the specifications defined in Table 128A–1 if measured at TP4_{H-D} (see Figure 128A–4).

Table 128A–1—2.5GSEI host output characteristics

Parameter	Subclause reference	Value	Units
Signaling rate	128A.3.1.1	3.125 ± 100 ppm	GBd
Nominal unit interval	128A.3.1.1	320	ps
DC common-mode output voltage (max)	128A.3.1.2	1.9	V
AC common-mode output voltage (max, RMS)	128A.3.1.2	30	mV
Peak-to-peak differential output voltage (min) transmitter enabled	128.7.1.4	580	mV
Peak-to-peak differential output voltage (max) transmitter enabled	128.7.1.4	1200	mV
Peak-to-peak differential output voltage (max) transmitter disabled	128.7.1.4	35	mV
Maximum transition time (20%–80%)	128.7.1.7	460	ps
Differential output return loss (min)	128A.3.1.3	See Equation (128A–2)	dB
Max output jitter (peak-to-peak)	128A.3.1.5	0.2	UI
Random jitter		0.12	UI
Deterministic jitter		0.035	UI
Duty Cycle Distortion ¹		0.32	UI
Total jitter			
Signal-to-noise-and-distortion-ratio (min) (SNDR)	128A.3.1.6	25	dB

¹Duty Cycle Distortion is considered part of deterministic jitter distribution.

A test system, as depicted in Figure 128B–1, with a fourth-order Bessel-Thomson low-pass response with 8 GHz 3 dB bandwidth, is to be used for all output signal measurements, unless otherwise specified.

128A.3.1.1 Signaling rate and range

The 2.5GSEI signaling rate is specified in Table 128A–1. The nominal unit interval is the inverse of the signaling rate.