

INTERNATIONAL
STANDARD

ISO/
IEC/IEEE
8802-3

Third edition
2021-02

AMENDMENT 13
2022-10

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

**Part 3:
Standard for Ethernet**

**AMENDMENT 13: Physical layers and
management parameters for 100 Gb/s
operation over DWDM systems**

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

Partie 3: Norme pour Ethernet

*AMENDEMENT 13: Couches physiques et paramètres de gestion pour
un fonctionnement à 100 Gb/s sur réseaux DWDM*



Reference number
ISO/IEC/IEEE 8802-3:2021/Amd.13:2022(E)

© IEEE 2021



COPYRIGHT PROTECTED DOCUMENT

© IEEE 2021

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. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives or www.iec.ch/members_experts/refdocs).

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 <https://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.13 was prepared by the LAN/MAN of the IEEE Computer Society (as IEEE 802.3ct-2021) 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 8802-3 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.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

IEEE Std 802.3ct™-2021

(Amendment to IEEE Std 802.3™-2018
as amended by IEEE Std 802.3cb™-2018,
IEEE Std 802.3bt™-2018,
IEEE Std 802.3cd™-2018,
IEEE Std 802.3cn™-2019,
IEEE Std 802.3cg™-2019,
IEEE Std 802.3cq™-2020,
IEEE Std 802.3cm™-2020,
IEEE Std 802.3ch™-2020,
IEEE Std 802.3ca™-2020,
IEEE Std 802.3cr™-2021,
IEEE Std 802.3cu™-2021,
and IEEE Std 802.3cv™-2021)

IEEE Standard for Ethernet

Amendment 13: Physical Layers and Management Parameters for 100 Gb/s Operation over DWDM Systems

Developed by the

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 16 June 2021

IEEE SA Standards Board

IECNORM.COM Click to view the full PDF of IEEE Std 802.3:2021/Amendment 13:2022

Abstract: This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 152 through Clause 154 and Annex 154A. This amendment adds 100 Gb/s Physical Layer specifications and management parameters for operation over DWDM systems using a combination of phase and amplitude modulation with coherent detection for reaches of at least 80 km.

Keywords: 100 Gb/s Ethernet, 100GBASE-ZR, DP-DQPSK, dense wavelength division multiplexing (DWDM), DWDM black link, Energy-Efficient Ethernet (EEE), Ethernet, forward error correction (FEC), IEEE 802.3™, IEEE 802.3ct™, Physical Medium Dependent (PMD) sublayer, single-mode fiber (SMF)

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

Copyright © 2021 by The Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 9 July 2021. 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 978-1-5044-7694-2 STD24792
Print: ISBN 978-1-5044-7695-9 STDPD24792

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 Standards documents are made available for use subject to important notices and legal disclaimers. These notices and disclaimers, or a reference to this page (<https://standards.ieee.org/ipr/disclaimers.html>), appear in all standards and may be found under the heading “Important Notices and Disclaimers Concerning IEEE Standards Documents.”

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

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE SA) Standards Board. IEEE develops its standards through an accredited consensus development process, which brings together volunteers representing varied viewpoints and interests to achieve the final product. IEEE Standards are documents developed by volunteers with scientific, academic, and industry-based expertise in technical working groups. Volunteers are not necessarily members of IEEE or IEEE SA, and participate without compensation from IEEE. While IEEE administers the process and establishes rules to promote fairness in the consensus development process, IEEE does not independently evaluate, test, or verify the accuracy of any of the information or the soundness of any judgments contained in its standards.

IEEE does not warrant or represent the accuracy or completeness 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. In addition, IEEE does not warrant or represent that the use of the material contained in its standards is free from patent infringement. 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: THE NEED TO PROCURE 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 is 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, nor 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 the presenter's views should be considered the personal views of that individual rather than the formal position of IEEE, IEEE SA, the Standards Committee, or the Working Group.

Comments on standards

Comments for revision of IEEE Standards documents are welcome from any interested party, regardless of membership affiliation with IEEE or IEEE SA. However, **IEEE does not provide interpretations, 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 evaluating comments or in revisions to an IEEE standard is welcome to join the relevant IEEE working group. You can indicate interest in a working group using the Interests tab in the Manage Profile & Interests area of the [IEEE SA myProject system](#). An IEEE Account is needed to access the application.

Comments on standards should be submitted using the [Contact Us](#) form.

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 constitute 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.

Data privacy

Users of IEEE Standards documents should evaluate the standards for considerations of data privacy and data ownership in the context of assessing and using the standards in compliance with applicable laws and regulations.

Copyrights

IEEE draft and approved standards are copyrighted by IEEE under US 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 licensing fees, 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; <https://www.copyright.com/>. 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 10 years. When a document is more than 10 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 [IEEE Xplore](#) or [contact IEEE](#). For more information about the IEEE SA or IEEE's standards development process, visit the IEEE SA Website.

Errata

Errata, if any, for all IEEE standards can be accessed on the [IEEE SA Website](#). Search for standard number and year of approval to access the web page of the published standard. Errata links are located under the Additional Resources Details section. Errata are also available in [IEEE Xplore](#). Users are encouraged to periodically check for errata.

Patents

IEEE Standards are developed in compliance with the [IEEE SA Patent Policy](#).

IMPORTANT NOTICE

IEEE Standards do not guarantee or ensure safety, security, health, or environmental protection, or ensure against interference with or from other devices or networks. IEEE Standards development activities consider research and information presented to the standards development group in developing any safety recommendations. Other information about safety practices, changes in technology or technology implementation, or impact by peripheral systems also may be pertinent to safety considerations during implementation of the standard. 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.

Participants

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

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

John D'Ambrosia, *IEEE P802.3ct 100 Gb/s over DWDM systems Task Force Chair*
Tom Issenhuth, *IEEE P802.3ct 100 Gb/s over DWDM systems Task Force Editor-in-Chief*
Editor for supporting clauses, Phase II
Pete Anslow, *IEEE P802.3ct 100 Gb/s over DWDM systems Task Force*
Editor for supporting clauses, Phase I
Steve Trowbridge, *IEEE P802.3ct 100 Gb/s over DWDM systems Task Force*
Editor for Clause 152, Clause 153, Annex 83C and Annex 135A
Peter Stassar, *IEEE P802.3ct 100 Gb/s over DWDM systems Task Force*
Editor for Clause 154

David Abramson	Takashi Fukuoka	Daniel Koppermueller
Michikazu Aono	Ali Ghiasi	Glen Kramer
Nobuyasu Araki	Joel Goergen	Taketo Kumada
Tim Baggett	Steven Gorshe	Hans Lackner
Thananya Baldwin	Hideki Goto	Frank Lambrecht
Steven Baumgartner	Steffen Graber	Mark Laubach
Denis Beaudoin	Olaf Grau	Greg Le Cheminant
Gitesh Bhagwat	Robert Grow	David Lewis
Rich Boyer	Martin Gubow	Mike-Peng Li
David Brandt	Mark Gustlin	Alex Lin
Ralf-Peter Braun	Marek Hajduczenia	Robert Lingle
Theodore Brillhart	Howard Heck	Hai-Feng Liu
Paul Brooks	David Hess	William Lo
Matthew Brown	Brian Holden	Yuanqiu Luo
Leon Bruckman	Bernd Hormmeyer	Kent Lusted
Jairo Bustos Heredia	Yasuhiro Hyakutake	Jeffery Maki
Adrian Butter	Jonathan Ingham	David Malicoat
John Calvin	Kazuhiko Ishibe	Eric Maniloff
Clark Carty	Hideki Isono	Flavio Marques
David Chalupsky	Hiroaki Ito	Arthur Marris
Jacky Chang	Kenneth Jackson	Takeo Masuda
Xin Chang	Andrew Jimenez	Mick McCarthy
Chan Chen	John Johnson	Brett McClellan
Golam Choudhury	Chad Jones	Larry McMillan
Keng Hua Chuang	Peter Jones	Greg McSorley
Piers J. G. Dawe	Lokesh Kabra	Richard Mellitz
Fred Dawson	Haysam Kadry	Shimon Muller
Gerrit den Besten	Manabu Kagami	Sean Murphy
Claudio DeSanti	Upen Karetí	James Nadolny
Curtis Donahue	Athanasios Kasapi	Edward Nakamoto
Liang Du	Yong Kim	Raymond Nering
Kathryn Dube	Mark Kimber	Paul Neveux
Mike Dudek	Michael Klempa	Gary Nicholl
Frank Effenberger	Curtis Knittle	Shawn Nicholl
David Estes	Elizabeth Kochuparambil	Kevin Noll
John Ewen	Sam Kocsis	Mark Nowell
Vincent Ferretti	Wojciech Koczwarra	David Ofelt
Brian Franchuk	Paul Kolesar	Ryo Okabe
Matthias Fritsche	Taiji Kondo	

Tom Palkert	Masato Shiino	Alexander Umnov
Carlos Pardo	Ramin Shirani	Prasad Venugopal
Earl Parsons	Kapil Shrikhande	Edward Walter
Gerald Pepper	Jeff Slavick	Roy Wang
Rubén Perez De Aranda Alonso	Scott Sommers	Xuehuan Wang
David Piehler	Massimo Sorbara	James Weaver
Fabio Pittala	Edward Sprague	Dong Wei
Christopher Pohl	Heath Stewart	Brian Welch
William Powell	Junqing Sun	Matthias Wendt
Rick Rabinovich	Steve Swanson	Natalie Wienckowski
Parthasarathy Raju	Tomoo Takahara	Dance Wu
Adee Ran	Satoshi Takahashi	Peter Wu
Alon Regev	Tadashi Takahashi	Dayin Xu
Duane Remein	Kazuya Takayama	Yu Xu
Victor Renteria	Michael Takefman	James Young
Thomas Rettig	Masaru Terada	Lennart Yseboodt
Toshiaki Sakai	Geoffrey Thompson	Conrad Zerna
Hamid Salehi	Pirooz Tooyserkani	Xingxin Zhang
Sam Sambasivan	Nathan Tracy	Chunhui Zhu
Edward Sayre	Viet Tran	Yan Zhuang
Matthew Schmitt	David Tremblay	Martin Zielinski
Hossein Sedarat	Mike Tu	George Zimmerman
Masood Shariff	Ed Ulrichs	Pavel Zivny

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

Robert Aiello	Werner Hoelzl	Adee Ran
Thomas Alexander	Thomas Huber	R. K. Rannow
Greg Armstrong	Gergely Huszak	Lakshman Raut
Rich Boyer	Yasuhiro Hyakutake	Robert Robinson
Ralf-Peter Braun	Jonathan Ingham	Benjamin Rolfe
Theodore Brillhart	Osamu Ishida	Toshiaki Sakai
Matthew Brown	Hideki Isono	Olindo Savi
Jairo Bustos Heredia	Tom Issenhuth	Matthew Schmitt
William Byrd	SangKwon Jeong	Qingya She
Steven B. Carlson	Jack Jewell	Scott Sommers
Juan Carreon	Peter Jones	Dorothy Stanley
Clark Carty	Lokesh Kabra	Peter Stassar
Pin Chang	Stuart Kerry	Walter Struppler
Chan Chen	Yong Kim	Mitsutoshi Sugawara
Charles Cook	Mark Laubach	Kohichi Tamura
John D'Ambrosia	Pi-Cheng Law	James Theodoras
Piers J. G. Dawe	David Lewis	Nathan Tracy
John Deandrea	Jon Lewis	David Tremblay
Christopher Diminico	Karen Liu	Stephen Trowbridge
Frank Effenberger	Valerie Maguire	Mark-Rene Uchida
Avraham Freedman	Jeffery Maki	Dmitri Varsanofiev
Matthias Fritsche	Richard Mellitz	Ruoxu Wang
Limin Geng	Nick S. A. Nikjoo	Xinyuan Wang
Ali Ghiasi	Mark Nowell	Lisa Ward
Zhigang Gong	Satoshi Obara	Keith Waters
Marek Hajduczenia	Tom Palkert	James Weaver
Xiang He	Carlos Pardo	James Young
Adam Healey	Bansi Patel	Yu Yuan
Marco Hernandez	Arumugam Paventhan	Oren Yuen
David Hess	David Piehler	Bo Zhang
Shui Heung	Patty Polpattana	

When the IEEE SA Standards Board approved this standard on 16 June 2021, it had the following membership:

Gary Hoffman, *Chair*
Jon Walter Rosdahl, *Vice Chair*
John D. Kulick, *Past Chair*
Konstantinos Karachalios, *Secretary*

Edward A. Addy
Doug Edwards
Ramy Ahmed Fathy
J. Travis Griffith
Thomas Koshy
Joseph L. Koepfinger*
David J. Law

Howard Li
Daozhuang Lin
Kevin Lu
Daleep C. Mohla
Chenhui Niu
Damir Novosel
Annette Reilly
Dorothy Stanley

Mehmet Ulema
Lei Wang
F. Keith Waters
Karl Weber
Sha Wei
Howard Wolfman
Daidi Zhong

*Member Emeritus

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

Introduction

This introduction is not part of IEEE Std 802.3ct-2021, IEEE Standard for Ethernet—Amendment 13: Physical Layers and Management Parameters for 100 Gb/s Operation over DWDM Systems.

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 IEEE Std 802.3ct-2021 publication, 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 include general information on 40 Gb/s and 100 Gb/s operation as well as 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 as 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.

IEEE Std 802.3bt™-2018

Amendment 2—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 145, Annex 145A, Annex 145B, and Annex 145C. This amendment adds power delivery using all four pairs in the structured wiring plant, resulting in greater power being available to end devices. This amendment also allows for lower standby power consumption in end devices and adds a mechanism to better manage the available power budget.

IEEE Std 802.3cd™-2018

Amendment 3—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 131 through Clause 140 and Annex 135A through Annex 136D. This amendment adds MAC parameters, Physical Layers, and management parameters for the transfer of IEEE 802.3 format frames at 50 Gb/s, 100 Gb/s, and 200 Gb/s.

IEEE Std 802.3cn™-2019

Amendment 4—This amendment includes changes to IEEE Std 802.3-2018 and adds 50 Gb/s, 200 Gb/s, and 400 Gb/s Physical Layer specifications and management parameters for operation over single-mode fiber with reaches of at least 40 km.

IEEE Std 802.3cg™-2019

Amendment 5—This amendment includes changes to IEEE Std 802.3-2018 and its amendments and adds Clause 146 through Clause 148 and Annex 146A and Annex 146B. This amendment adds 10 Mb/s Physical Layer specifications and management parameters for operation on a single balanced pair of conductors.

IEEE Std 802.3cq™-2020

Amendment 6—This amendment includes editorial and technical corrections, refinements, and clarifications to Clause 33 and related portions of the standard.

IEEE Std 802.3cm™-2020

Amendment 7—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 150. This amendment adds Physical Layer (PHY) specifications and management parameters for 400 Gb/s operation on four pairs (400GBASE-SR4.2) and eight pairs (400GBASE-SR8) of multimode fiber, over reaches of at least 100 m.

IEEE Std 802.3ch™-2020

Amendment 8—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 149, Annex 149A, Annex 149B, and Annex 149C. This amendment adds physical layer specifications and management parameters for operation at 2.5 Gb/s, 5 Gb/s, and 10 Gb/s over a single balanced pair of conductors.

IEEE Std 802.3ca™-2020

Amendment 9—This amendment to IEEE Std 802.3-2018 extends the operation of Ethernet passive optical networks (EPONs) to multiple channels of 25 Gb/s providing both symmetric and asymmetric operation for the following data rates (downstream/upstream): 25/10 Gb/s, 25/25 Gb/s, 50/10 Gb/s, 50/25 Gb/s, and 50/50 Gb/s. This amendment specifies the 25 Gb/s EPON Multi-Channel Reconciliation Sublayer (MCRS), Nx25G-EPON Physical Coding Sublayers (PCSs), Physical Media Attachment (PMA) sublayers, and Physical Medium Dependent (PMD) sublayers that support both symmetric and asymmetric data rates while maintaining backward compatibility with already deployed 10 Gb/s EPON equipment. The EPON operation is defined for distances of at least 20 km, and for a split ratio of at least 1:32.

IEEE Std 802.3cr™-2021

Amendment 10—This amendment includes changes to IEEE Std 802.3-2018 and adds Annex J. This amendment replaces references to the IEC 60950 series of standards (including IEC 60950-1 “Information technology equipment—Safety—Part 1: General requirements”) with appropriate references to the IEC 62368 “Audio/video, information and communication technology equipment” series and makes appropriate changes to the standard corresponding to the new references.

IEEE Std 802.3cu™-2021

Amendment 11—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 151. This amendment adds Physical Layer (PHY) specifications and management parameters for 100 Gb/s and 400 Gb/s operation over single-mode fiber, based on 100 Gb/s per wavelength optical signaling.

IEEE Std 802.3cv™-2021

Amendment 12—This amendment includes editorial and technical corrections, refinements, and clarifications to Clause 145, Power over Ethernet, and related portions of the standard.

IEEE Std 802.3ct™-2021

Amendment 13—This amendment includes changes to IEEE Std 802.3-2018 and adds Clause 152 through Clause 154 and Annex 154A. This amendment adds 100 Gb/s Physical Layer specifications and management parameters for operation over DWDM systems with reaches of at least 80 km.

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.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

Contents

1. Introduction.....	21
1.3 Normative references	21
1.4 Definitions	21
1.5 Abbreviations.....	22
30. Management.....	23
30.2 Managed Objects	23
30.2.5 Capabilities.....	23
30.5 Layer management for medium attachment units (MAUs).....	23
30.5.1 MAU managed object class.....	23
30.5.1.1 MAU attributes	23
30.5.1.1.2 aMAUType	23
30.5.1.1.15 aFECAbility	24
30.5.1.1.16 aFECmode.....	24
30.5.1.1.17 aFECCorrectedBlocks.....	24
30.5.1.1.18 aFECUncorrectableBlocks.....	25
30.5.1.1.26 aRSFECBIPErrCount	25
30.5.1.1.27 aRSFECLaneMapping	26
30.5.1.1.27a aSCFECLaneMapping	26
30.5.1.1.28 aRSFECBypassAbility.....	26
30.5.1.1.29 aRSFECBypassIndicationAbility.....	27
30.5.1.1.32 aPCSFECBypassIndicationAbility	27
45. Management Data Input/Output (MDIO) Interface.....	28
45.2 MDIO Interface Registers.....	28
45.2.1 PMA/PMD registers.....	28
45.2.1.6 PMA/PMD control 2 register (Register 1.7)	29
45.2.1.7 PMA/PMD status 2 register (Register 1.8).....	29
45.2.1.7.4 Transmit fault (1.8.11)	29
45.2.1.7.5 Receive fault (1.8.10).....	30
45.2.1.8 PMD transmit disable register (Register 1.9).....	30
45.2.1.21b 40G/100G PMA/PMD extended ability 2 register (Register 1.26)	30
45.2.1.21b.3aa 100GBASE-ZR ability (1.26.6)	31
45.2.1.133a Tx optical channel control register (Register 1.800)	31
45.2.1.133a.1 Tx optical channel index (1.800.5:0)	31
45.2.1.133b Tx optical channel ability 1 register (Register 1.801)	31
45.2.1.133b.1 Tx index ability 0 through 15 (1.801.0 through 1.801.15).....	32
45.2.1.133c Tx optical channel ability 2 register (Register 1.802)	32
45.2.1.133c.1 Tx index ability 16 through 31 (1.802.0 through 1.802.15).....	33
45.2.1.133d Tx optical channel ability 3 register (Register 1.803)	34
45.2.1.133d.1 Tx index ability 32 through 47 (1.803.0 through 1.803.15).....	34
45.2.1.133e Rx optical channel control register (Register 1.820).....	35
45.2.1.133e.1 Tx Rx different optical channel ability (1.820.15).....	35
45.2.1.133e.2 Rx optical channel index (1.820.5:0).....	35
45.2.1.133f Rx optical channel ability 1 register (Register 1.821)	36
45.2.1.133f.1 Rx index ability 0 through 15 (1.821.0 through 1.821.15)	36
45.2.1.133g Rx optical channel ability 2 register (Register 1.822).....	37
45.2.1.133g.1 Rx index ability 16 through 31 (1.822.0 through 1.822.15)	38

45.2.1.133h	Rx optical channel ability 3 register (Register 1.823)	38
45.2.1.133h.1	Rx index ability 32 through 47 (1.823.0 through 1.823.15)	39
45.2.1.186aa	IFEC control register (Register 1.2200)	39
45.2.1.186aa.1	IFEC bypass indication enable (1.2200.1)	39
45.2.1.186aa.2	IFEC bypass correction enable (1.2200.0)	39
45.2.1.186ab	IFEC status register (Register 1.2201)	39
45.2.1.186ab.1	PCS align status (1.2201.15)	40
45.2.1.186ab.2	IFEC align status (1.2201.14)	40
45.2.1.186ab.3	IFEC AM lock 3 (1.2201.11)	40
45.2.1.186ab.4	IFEC AM lock 2 (1.2201.10)	41
45.2.1.186ab.5	IFEC AM lock 1 (1.2201.9)	41
45.2.1.186ab.6	IFEC AM lock 0 (1.2201.8)	41
45.2.1.186ab.7	IFEC high SER (1.2201.2)	41
45.2.1.186ab.8	IFEC bypass indication ability (1.2201.1)	41
45.2.1.186ab.9	IFEC bypass correction ability (1.2201.0)	41
45.2.1.186ac	IFEC corrected codewords counter (Register 1.2202, 1.2203)	41
45.2.1.186ad	IFEC uncorrected codewords counter (Register 1.2204, 1.2205)	42
45.2.1.186ae	IFEC lane mapping register (Register 1.2206)	42
45.2.1.186af	IFEC symbol error counter, lane 0 (Register 1.2210, 1.2211)	43
45.2.1.186ag	IFEC symbol error counter, lane 1 through 3 (Register 1.2212, 1.2213, 1.2214, 1.2215, 1.2216, 1.2217)	43
45.2.1.186ah	SC-FEC alignment status 1 register (Register 1.2246)	43
45.2.1.186ah.1	SC-FEC align status (1.2246.12)	43
45.2.1.186ah.2	SC-FEC FAS lock 7 (1.2246.7)	44
45.2.1.186ah.3	SC-FEC FAS lock 6 (1.2246.6)	44
45.2.1.186ah.4	SC-FEC FAS lock 5 (1.2246.5)	44
45.2.1.186ah.5	SC-FEC FAS lock 4 (1.2246.4)	45
45.2.1.186ah.6	SC-FEC FAS lock 3 (1.2246.3)	45
45.2.1.186ah.7	SC-FEC FAS lock 2 (1.2246.2)	45
45.2.1.186ah.8	SC-FEC FAS lock 1 (1.2246.1)	45
45.2.1.186ah.9	SC-FEC FAS lock 0 (1.2246.0)	45
45.2.1.186ai	SC-FEC alignment status 2 register (Register 1.2247)	45
45.2.1.186ai.1	SC-FEC FAS lock 19 (1.2247.11)	46
45.2.1.186ai.2	SC-FEC FAS lock 18 (1.2247.10)	46
45.2.1.186ai.3	SC-FEC FAS lock 17 (1.2247.9)	46
45.2.1.186ai.4	SC-FEC FAS lock 16 (1.2247.8)	47
45.2.1.186ai.5	SC-FEC FAS lock 15 (1.2247.7)	47
45.2.1.186ai.6	SC-FEC FAS lock 14 (1.2247.6)	47
45.2.1.186ai.7	SC-FEC FAS lock 13 (1.2247.5)	47
45.2.1.186ai.8	SC-FEC FAS lock 12 (1.2247.4)	47
45.2.1.186ai.9	SC-FEC FAS lock 11 (1.2247.3)	47
45.2.1.186ai.10	SC-FEC FAS lock 10 (1.2247.2)	47
45.2.1.186ai.11	SC-FEC FAS lock 9 (1.2247.1)	47
45.2.1.186ai.12	SC-FEC FAS lock 8 (1.2247.0)	48
45.2.1.186aj	SC-FEC lane mapping, lane 0 register (Register 1.2250)	48
45.2.1.186ak	SC-FEC lane mapping, lane 1 through 19 registers (Registers 1.2251 through 1.2269)	48
45.2.1.186al	SC-FEC corrected codewords counter (Register 1.2276, 1.2277)	48
45.2.1.186am	SC-FEC uncorrected codewords counter (Register 1.2278, 1.2279)	49
45.2.1.186an	SC-FEC total bits register (Register 1.2280, 1.2281, 1.2282, 1.2283)	49
45.2.1.186ao	SC-FEC corrected bits register (Register 1.2284, 1.2285, 1.2286, 1.2287)	49
78.	Energy-Efficient Ethernet (EEE)	51

78.1	Overview.....	51
78.1.4	PHY types optionally supporting EEE.....	51
80.	Introduction to 40 Gb/s and 100 Gb/s networks.....	52
80.1	Overview.....	52
80.1.3	Relationship of 40 Gigabit and 100 Gigabit Ethernet to the ISO OSI reference model.....	52
80.1.4	Nomenclature.....	52
80.1.5	Physical Layer signaling systems.....	53
80.2	Summary of 40 Gigabit and 100 Gigabit Ethernet sublayers.....	53
80.2.2	Physical Coding Sublayer (PCS).....	53
80.2.3	Forward Error Correction (FEC) sublayers.....	54
80.2.4	Physical Medium Attachment (PMA) sublayer.....	54
80.2.5	Physical Medium Dependent (PMD) sublayer.....	54
80.3	Service interface specification method and notation.....	54
80.3.2	Instances of the Inter-sublayer service interface.....	54
80.4	Delay constraints.....	56
80.5	Skew constraints.....	56
80.7	Protocol implementation conformance statement (PICS) proforma.....	57
82.	Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R.....	58
82.2	Physical Coding Sublayer (PCS).....	58
82.2.3	64B/66B transmission code.....	58
82.2.3.3	Block structure.....	58
152.	Inverse RS-FEC sublayer.....	59
152.1	Overview.....	59
152.1.1	Scope.....	59
152.1.2	Position of Inverse RS-FEC in the 100GBASE-R sublayers.....	59
152.2	Inverse RS-FEC service interface.....	59
152.3	PMA or FEC sublayer compatibility.....	59
152.4	Delay constraints.....	59
152.5	Functions within the Inverse RS-FEC sublayer.....	60
152.5.1	Functional block diagram.....	60
152.5.2	Transmit function.....	61
152.5.2.1	Alignment lock and deskew.....	61
152.5.2.2	Lane reorder.....	62
152.5.2.3	Reed-Solomon decoder.....	62
152.5.2.4	Alignment marker removal.....	62
152.5.2.5	256B/257B to 64B/66B transcoder.....	62
152.5.2.6	Block distribution.....	63
152.5.2.7	Alignment marker mapping and insertion.....	63
152.5.2.8	Transmit bit ordering.....	64
152.5.3	Receive function.....	66
152.5.3.1	Lane block synchronization.....	66
152.5.3.2	Alignment lock and deskew.....	66
152.5.3.3	Lane reorder.....	66
152.5.3.4	Alignment marker removal.....	66
152.5.3.5	64B/66B to 256B/257B transcoder.....	66
152.5.3.6	Alignment marker mapping and insertion.....	67
152.5.3.7	Reed-Solomon encoder.....	69

152.5.3.8	Symbol distribution	69
152.5.3.9	Receive bit ordering.....	69
152.5.4	Detailed functions and state diagrams.....	71
152.5.4.1	State diagram conventions.....	71
152.5.4.2	State variables.....	71
152.5.4.2.1	Variables	71
152.5.4.2.2	Functions.....	72
152.5.4.2.3	Counters	73
152.5.4.3	State diagrams.....	73
152.6	Inverse RS-FEC MDIO function mapping.....	73
152.6.1	IFEC_bypass_correction_enable.....	73
152.6.2	IFEC_bypass_indication_enable.....	74
152.6.3	IFEC_bypass_correction_ability.....	74
152.6.4	IFEC_bypass_indication_ability.....	75
152.6.5	hi_ser.....	75
152.6.6	amps_lock<x>	75
152.6.7	IFEC_align_status	75
152.6.8	IFEC_corrected_cw_counter.....	75
152.6.9	IFEC_uncorrected_cw_counter.....	75
152.6.10	IFEC_lane_mapping<x>	75
152.6.11	IFEC_symbol_error_counter_i.....	76
152.6.12	align_status.....	76
152.6.13	BIP_error_counter_i.....	76
152.6.14	lane_mapping<x>.....	76
152.6.15	block_lock<x>.....	76
152.6.16	am_lock<x>.....	76
152.7	Protocol implementation conformance statement (PICS) proforma for Clause 152, Inverse RS-FEC sublayer.....	77
152.7.1	Introduction	77
152.7.2	Identification	77
152.7.2.1	Implementation identification.....	77
152.7.2.2	Protocol summary.....	77
152.7.3	Major capabilities/options	78
152.7.4	PICS proforma tables for Inverse RS-FEC sublayer.....	78
152.7.4.1	Transmit function	78
152.7.4.2	Receive Function	79
152.7.4.3	State diagrams.....	80
153.	SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs.....	81
153.1	Overview.....	81
153.1.1	Scope	81
153.1.2	Position of SC-FEC and 100GBASE-ZR PMA in the 100GBASE-R sublayers.....	81
153.2	SC-FEC sublayer	81
153.2.1	FEC service interface	81
153.2.2	Delay constraints	82
153.2.3	Functions within the SC-FEC sublayer.....	83
153.2.3.1	Functional block diagram	83
153.2.3.2	Transmit function	83
153.2.3.2.1	Lane block synchronization	83
153.2.3.2.2	Alignment lock and deskew.....	84
153.2.3.2.3	Lane reorder	84
153.2.3.2.4	GMP mapper.....	84

153.2.3.2.5	SC-FEC encoder	87
153.2.3.2.6	Scrambler	87
153.2.3.2.7	Lane distribution	88
153.2.3.3	Receive function	89
153.2.3.3.1	Lane lock and deskew	89
153.2.3.3.2	Lane reorder	89
153.2.3.3.3	De-scrambler	89
153.2.3.3.4	SC-FEC decoder	89
153.2.3.3.5	GMP demapper	89
153.2.3.3.6	Block alignment	90
153.2.3.3.7	Block distribution	90
153.2.4	Detailed functions and state diagrams	90
153.2.4.1	State variables	90
153.2.4.1.1	Variables	90
153.2.4.2	Functions	91
153.2.4.3	Counters	92
153.2.4.4	State diagrams	92
153.2.5	SC-FEC MDIO function mapping	94
153.2.5.1	FEC_corrected_cw_counter	94
153.2.5.2	FEC_uncorrected_cw_counter	94
153.2.5.3	FEC_total_bits_counter	94
153.2.5.4	FEC_corrected_bits_counter	95
153.3	100GBASE-ZR PMA sublayer	95
153.3.1	100GBASE-ZR PMA service interface	95
153.3.2	Functions within the 100GBASE-ZR PMA sublayer	96
153.3.2.1	Functional block diagram	96
153.3.2.2	Transmit function	96
153.3.2.2.1	Lane interleave	96
153.3.2.2.2	DQPSK encode	96
153.3.2.3	Receive function	97
153.3.2.3.1	DQPSK decode	97
153.3.2.3.2	Lane de-interleave	97
153.4	Protocol implementation conformance statement (PICS) proforma for Clause 153, SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs	98
153.4.1	Introduction	98
153.4.2	Identification	98
153.4.2.1	Implementation identification	98
153.4.2.2	Protocol summary	98
153.4.3	Major capabilities/options	99
153.4.4	PICS proforma tables for SC-FEC sublayer for 100GBASE-ZR PHYs	99
153.4.4.1	Transmit function	99
153.4.4.2	Receive function	100
153.4.4.3	State diagrams	100
154.	Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-ZR	101
154.1	Overview	101
154.1.1	Bit error ratio	102
154.2	Physical Medium Dependent (PMD) service interface	103
154.3	Delay and Skew	103
154.3.1	Delay constraints	103
154.3.2	Skew constraints	103
154.4	PMD MDIO function mapping	104

154.5	PMD functional specifications.....	105
154.5.1	PMD block diagram	105
154.5.2	PMD transmit function.....	105
154.5.3	PMD receive function	106
154.5.4	PMD global signal detect function.....	106
154.5.5	PMD reset function	106
154.5.6	PMD global transmit disable function (optional).....	106
154.5.7	PMD fault function (optional).....	106
154.5.8	PMD transmit fault function (optional).....	107
154.5.9	PMD receive fault function (optional)	107
154.6	DWDM channel over a DWDM black link	107
154.7	PMD to MDI optical specifications for 100GBASE-ZR.....	109
154.7.1	100GBASE-ZR transmitter optical specifications	110
154.7.2	100GBASE-ZR receive optical specifications.....	111
154.8	100GBASE-ZR DWDM black link transfer characteristics.....	111
154.9	Definition of optical parameters and measurement methods.....	111
154.9.1	Test patterns for optical parameters	111
154.9.2	Optical center frequency (wavelength) and side-mode suppression ratio (SMSR)	113
154.9.3	Average channel output power.....	113
154.9.4	Spectral excursion	113
154.9.5	Laser linewidth.....	113
154.9.6	Offset between the carrier and the nominal center frequency.....	114
154.9.7	Power difference between X and Y polarizations	114
154.9.8	Skew between X and Y polarizations	114
154.9.9	Error vector magnitude.....	114
154.9.10	I-Q offset	114
154.9.11	Optical signal-to-noise ratio (OSNR).....	114
154.9.12	Transmitter in-band OSNR	114
154.9.13	Average receive power.....	114
154.9.14	Receiver sensitivity	115
154.9.15	Receiver OSNR	115
154.9.16	Receiver OSNR tolerance	115
154.9.17	Ripple	115
154.9.18	Optical path OSNR penalty.....	115
154.9.19	Optical path power penalty	115
154.9.20	Polarization dependent loss.....	115
154.9.21	Polarization rotation speed.....	116
154.9.22	Inter-channel crosstalk at TP3.....	116
154.9.23	Interferometric crosstalk at TP3.....	116
154.10	Safety, installation, environment, and labeling.....	116
154.10.1	General safety.....	116
154.10.2	Laser safety	116
154.10.3	Installation.....	116
154.10.4	Environment.....	116
154.10.5	Electromagnetic emission	117
154.10.6	Temperature, humidity, and handling	117
154.10.7	PMD labeling requirements	117
154.11	Medium Dependent Interface (MDI).....	117
154.12	Protocol implementation conformance statement (PICS) proforma for Clause 154, Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-ZR.....	118
154.12.1	Introduction	118
154.12.2	Identification	118
154.12.2.1	Implementation identification.....	118
154.12.2.2	Protocol summary.....	118

154.12.3	Major capabilities/options	119
154.12.4	PICS proforma tables for Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-ZR	119
154.12.4.1	PMD functional specifications	119
154.12.4.2	Management functions	120
154.12.4.3	PMD to MDI optical specifications for 100GBASE-ZR.....	120
154.12.4.4	Optical measurement methods.....	121
154.12.4.5	Environmental specifications	121
154.12.4.6	Characteristics of DWDM black link and MDI.....	121
Annex A (informative) Bibliography		122
Annex 83C (informative) PMA sublayer partitioning examples.....		123
83C.4	Partitioning examples with SC-FEC.....	123
83C.4.1	CAUI-4 with SC-FEC.....	123
Annex 135A (informative) 50 Gb/s and 100 Gb/s PMA sublayer partitioning examples.....		124
135A.3	Partitioning examples of 100GAUI-n with Inverse RS-FEC	124
135A.3.1	100GAUI-n with Inverse RS-FEC.....	124
135A.3.2	CAUI-4 chip-to-chip and 100GAUI-n chip-to-module with Inverse RS-FEC.....	125
Annex 154A (informative) Examples of 100GBASE-ZR compliant DWDM black links		126
154A.1	Introduction.....	126
154A.2	Relationship between OSNR and average optical power	127
154A.3	Examples of DWDM black link applications with OSNR at TP3 between 19.5 dB (12.5 GHz) and 35 dB (12.5 GHz).....	128
154A.4	Example of DWDM black link applications with OSNR at TP3 greater than or equal to 35 dB (12.5 GHz)	129

IEEE Standard for Ethernet

Amendment 13: Physical Layers and Management Parameters for 100 Gb/s Operation over DWDM Systems

(This amendment is based on IEEE Std 802.3™-2018, as amended by IEEE Std 802.3cb™-2018, IEEE Std 802.3bt™-2018, IEEE Std 802.3cd™-2018, IEEE Std 802.3cn™-2019, IEEE Std 802.3cg™-2019, IEEE Std 802.3cq™-2020, IEEE Std 802.3cm™-2020, IEEE Std 802.3ch™-2020, IEEE Std 802.3ca™-2020, IEEE Std 802.3cr™-2021, IEEE Std 802.3cu™-2021, and IEEE Std 802.3cv™-2021.)

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

Cross references that refer to clauses, tables, equations, or figures not covered by this amendment are highlighted in green.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 three references into 1.3 in alphanumeric order:

ITU-T Recommendation G.698.2—Amplified multichannel dense wavelength division multiplexing applications with single channel optical interfaces.²

ITU-T Recommendation G.709—Interfaces for the optical transport network.

ITU-T Recommendation G.709.2—OTU4 long-reach interface.

1.4 Definitions

Insert the following two new definitions after 1.4.35 “100GBASE-SR4”:

1.4.35a 100GBASE-Z: An IEEE 802.3 family of Physical Layer devices using 100GBASE-R encoding, a combination of phase and amplitude modulation, and coherent detection. (See IEEE Std 802.3, Clause 154.)

1.4.35b 100GBASE-ZR: IEEE 802.3 Physical Layer specification for a 100 Gb/s dense wavelength division multiplexing (DWDM) PHY using 100GBASE-R encoding, dual polarization differential quadrature phase shift keying (DP-DQPSK) modulation, and coherent detection with reach up to at least 80 km. (See IEEE Std 802.3, Clause 154.)

Insert the following new definition after 1.4.160 “bit time (BT)”:

1.4.160a black link approach: The specification of the input, output, and transfer characteristics of the unidirectional transmission path from TP2 to TP3 for a given dense wavelength division multiplexing (DWDM) channel within a DWDM black link, without specifying how the transmission path is implemented. (See, for example, IEEE Std 802.3, Clause 154, Figure 154–3.)

Insert the following new definition after 1.4.181 “Channel Operating Margin (COM)”:

1.4.181a channel spacing: The center-to-center difference in frequency or wavelength between adjacent channels in a WDM application. Dense wavelength division multiplexing (DWDM) channel spacings are based on the grid found in ITU-T G.694.1.

Insert the following new definition after 1.4.227 “defect”:

1.4.227a dense wavelength division multiplexing: An optical WDM technology where the frequency spacing is less than or equal to 1000 GHz.

Insert the following three new definitions after 1.4.237 “duplex channel”:

1.4.237a DWDM black link: An aggregate of pairs of dense wavelength division multiplexing (DWDM) channels, with each pair supporting one full duplex connection where the implementation of the transmission paths is not specified.

²ITU-T publications are available from the International Telecommunications Union (<https://www.itu.int/>).

1.4.237b DWDM channel: The transmission path from a transmitting DWDM PHY (TP2) to a receiving DWDM PHY (TP3).

1.4.237c DWDM PHY: An Ethernet PHY that transmits and receives on selected dense wavelength division multiplexing (DWDM) center frequencies for transmission over one selected DWDM channel in each direction.

Insert the following new definition after 1.4.400 “Point-to-point emulation (P2PE)” (renumbered from 1.4.401 due to the deletion of 1.4.294 by IEEE Std 802.3bt-2018):

1.4.400a polarization dependent loss: The variation of insertion loss due to a variation of the state of polarization of an optical signal over all states of polarization within the channel frequency or wavelength range.

1.5 Abbreviations

Insert the following new abbreviations into the list in 1.5, in alphanumeric order:

DP-DQPSK	dual polarization differential quadrature phase shift keying
DQPSK	differential quadrature phase shift keying
DWDM	dense wavelength division multiplexing
FAS	frame alignment signal
GMP	generic mapping procedure (see ITU-T G.709)
IFEC	inverse RS-FEC
MFAS	multi-frame alignment signal
OSNR	optical signal-to-noise ratio
SC-FEC	staircase FEC

30. Management

30.2 Managed Objects

30.2.5 Capabilities

Insert a new row for aSCFECLaneMapping in Table 30–1e after the row for “aRSFECLaneMapping” as follows (unchanged rows not shown):

Table 30–1e—Capabilities

			DTE	Repeater	MAU
			Basic Package (mandatory)		
			Mandatory Package (mandatory)		
			Recommended Package (optional)		
			Optional Package (optional)		
			Array Package (optional)		
			Excessive Deferral Package (optional)		
			Multiple PHY Package (optional)		
			PHY Error Monitor Capability (optional)		
			Basic Control Capability (mandatory)		
			Performance Monitor Capability (optional)		
			Address Tracking Capability (optional)		
			100/1000 Mb/s Monitor Capability (optional)		
			1000 Mb/s Burst Monitor Capability (optional)		
			Basic Package (mandatory)		
			MAU Control Package (optional)		
			Media Loss Tracking Package (conditional)		
			Broadband DTE MAU Package (conditional)		
			MII Capability (conditional)		
			PHY Error Monitor Capability (optional)		
			MultiGBASE-T Operating Margin package (conditional)		
			Forward Error Correction Package (conditional)		
			Energy-Efficient Ethernet (optional)		
			Auto-Negotiation Package (mandatory)		
oMAU managed object class (30.5.1)					
...					
aSCFECLaneMapping	ATTRIBUTE	GET			X
...					

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 100GBASE-ZR PHY type into the “APPROPRIATE SYNTAX” section of 30.5.1.1.2 after 100GBASE-ER4 (as inserted by IEEE Std 802.3cn-2019) as follows:

APPROPRIATE SYNTAX:

...	
100GBASE-ZR	100GBASE-R PCS/100GBASE-ZR PMA over a PMD with reach up to at least 80 km as specified in Clause 154
...	

30.5.1.1.15 aFECAbility

Change 30.5.1.1.15 as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

An ENUMERATED VALUE that meets the requirement of the description below

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

BEHAVIOUR DEFINED AS:

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

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC capability register (see 45.2.10.2 or 45.2.1.101).;

30.5.1.1.16 aFECmode

Change the first paragraph of the “BEHAVIOUR DEFINED AS” section of 30.5.1.1.16 as follows:

BEHAVIOUR DEFINED AS:

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

30.5.1.1.17 aFECCorrectedBlocks

Change 30.5.1.1.17 (as modified by IEEE Std 802.3cd-2018) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresettable counters. Each counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, 5 000 000 counts per second for 10 Gb/s, 25 Gb/s, and 40 Gb/s implementations, 10 000 000 counts per second for 50 Gb/s implementations, 2 500 000 counts per second for 100 Gb/s implementations, 40 000 000 counts per second for 200 Gb/s implementations, and 80 000 000 counts per second for 400 Gb/s implementations.

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/50/100/200/400GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs that support FEC across the MDI, an array of corrected FEC block counters. The counters do not increment for other PHY types. The indices of this array (0 to N – 1) denote the FEC sublayer instance number where N is the number of FEC sublayer instances in use. The number of FEC sublayer instances in use is set to one for PHYs that do not use PCS lanes or use a single FEC instance for all lanes. Each element of this array contains a count of corrected FEC blocks for that FEC sublayer instance.

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

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC corrected blocks

counter(s) (see 45.2.10.5 and 45.2.1.103 for 10GBASE-R, 45.2.3.41 for 10GBASE-PR and 10/1GBASE-PRX, 45.2.1.125 for BASE-R, 45.2.1.112 for RS-FEC, and 45.2.3.61 for PCS FEC, and 45.2.1.186al for SC-FEC).;

30.5.1.1.18 aFECUncorrectableBlocks

Change 30.5.1.1.18 (as modified by IEEE Std 802.3cd-2018) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresettable counters. Each counter has a maximum increment rate of 1 200 000 counts per second for 1000 Mb/s implementations, and 5 000 000 counts per second for 10 Gb/s, 25 Gb/s, and 40 Gb/s implementations, 10 000 000 counts per second for 50 Gb/s implementations, 2 500 000 counts per second for 100 Gb/s implementations, 40 000 000 counts per second for 200 Gb/s implementations, and 80 000 000 counts per second for 400 Gb/s implementations.

BEHAVIOUR DEFINED AS:

For 1000BASE-PX, 10/25/40/50/100/200/400GBASE-R, 100GBASE-P, 10GBASE-PR, or 10/1GBASE-PRX PHYs that support FEC across the MDI, an array of uncorrectable FEC block counters. The counters do not increment for other PHY types. The indices of this array (0 to N – 1) denote the FEC sublayer instance number where N is the number of FEC sublayer instances in use. The number of FEC sublayer instances in use is set to one for PHYs that do not use PCS lanes or use a single FEC instance for all lanes. Each element of this array contains a count of uncorrectable FEC blocks for that FEC sublayer instance.

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

If a Clause 45 MDIO Interface is present, then this attribute maps to the FEC uncorrectable blocks counter(s) (see 45.2.10.6 and 45.2.1.104 for 10GBASE-R, 45.2.3.42 for 10GBASE-PR and 10/1GBASE-PRX, 45.2.1.133 for BASE-R, 45.2.1.113 for RS-FEC, and 45.2.3.62 for PCS FEC, and 45.2.1.186am for SC-FEC).;

30.5.1.1.26 aRSFECBIPErrorCount

Change 30.5.1.1.26 (as modified by IEEE Std 802.3cd-2018) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of generalized nonresettable counters. Each counter has a maximum increment rate of 5 000 counts per second for 50 Gb/s and 100 Gb/s implementations.

BEHAVIOUR DEFINED AS:

For 50GBASE-R, 100GBASE-R, and 100GBASE-P PHYs that support RS-FEC across the MDI, an array of BIP error counters. The counters do not increment for other PHY types. The indices of this array (0 to N – 1) denote the PCS lane number where N is the number of PCS lanes in use. Each element of this array contains a count of BIP errors for that PCS lane.

Increment the counter by one for each BIP error across the MDI detected during alignment marker removal in the PCS for the corresponding lane.

If a Clause 45 MDIO Interface is present, then this attribute maps to the BIP error counters (see 45.2.1.117 and 45.2.1.118).;

30.5.1.1.27 aRSFECLaneMapping

Change 30.5.1.1.27 (as modified by IEEE Std 802.3cd-2018) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of INTEGERS.

BEHAVIOUR DEFINED AS:

For 50GBASE-R, 100GBASE-R, and 100GBASE-P PHYs that support RS-FEC across the MDI, an array of PCS lane identifiers. The indices of this array (0 to N – 1) denote the service interface lane number where N is the number of PCS lanes in use. Each element of this array contains the PCS lane number for the PCS lane across the MDI that has been detected in the corresponding service interface lane.

If a Clause 45 MDIO Interface is present, then this attribute maps to the Lane mapping registers (see 45.2.1.119 and 45.2.1.120).;

Insert new subclause 30.5.1.1.27a after 30.5.1.1.27 as follows:

30.5.1.1.27a aSCFECLaneMapping

ATTRIBUTE

APPROPRIATE SYNTAX:

A SEQUENCE of INTEGERS.

BEHAVIOUR DEFINED AS:

For a 100GBASE-R PHY that supports SC-FEC (see Clause 153) across the MDI, an array of PCS lane identifiers. The indices of this array (0 to N – 1) denote the service interface lane number where N is the number of PCS lanes in use. Each element of this array contains the PCS lane number for the PCS lane across the MDI that has been detected in the corresponding service interface lane.

If a Clause 45 MDIO Interface is present, then this attribute maps to the Lane mapping registers (see 45.2.1.186aj and 45.2.1.186ak).;

30.5.1.1.28 aRSFECBypassAbility

Change 30.5.1.1.28 as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

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

unknown	initializing, true state not yet known
supported	FEC bypass ability supported
not supported	FEC bypass ability not supported

BEHAVIOUR DEFINED AS:

A read-only value that indicates if ~~the a~~ PHY that supports RS-FEC across the MDI supports ~~an~~ the optional RS-FEC bypass ability (see 91.5.3.3).

For a PHY that does not support RS-FEC across the MDI, this attribute is not applicable.

If a Clause 45 MDIO Interface is present, then this attribute maps to the RS-FEC status register (see 45.2.1.111).;

30.5.1.1.29 aRSFECBypassIndicationAbility

Change 30.5.1.1.29 (as modified by IEEE Std 802.3cd-2018) as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

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

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

BEHAVIOUR DEFINED AS:

A read-only value that indicates if ~~the a PHY that supports RS-FEC across the MDI supports an~~ the optional RS-FEC error indication bypass ability.

For a PHY that does not support RS-FEC across the MDI, this attribute is not applicable.

If a Clause 45 MDIO Interface is present, then this attribute maps to the RS-FEC status register (see 45.2.1.111).;

30.5.1.1.32 aPCSFECBypassIndicationAbility

Change 30.5.1.1.32 as follows:

ATTRIBUTE

APPROPRIATE SYNTAX:

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

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

BEHAVIOUR DEFINED AS:

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

For a PHY that does not support RS-FEC across the MDI, this attribute is not applicable.

If a Clause 45 MDIO Interface is present, then this attribute maps to the PCS FEC status register (see 45.2.3.60).;

45. Management Data Input/Output (MDIO) Interface

45.2 MDIO Interface Registers

45.2.1 PMA/PMD registers

Change the reserved rows for 1.740 through 1.899 and 1.203 through 1.2293 in Table 45–3 (as modified by IEEE Std 802.3cg-2019) as follows (unchanged rows not shown):

Table 45–3—PMA/PMD registers

Register address	Register name	Subclause
...		
1.740 through 1.899 1.799	Reserved	
1.800	<u>Tx optical channel control</u>	45.2.1.133a
1.801	<u>Tx optical channel ability 1</u>	45.2.1.133b
1.802	<u>Tx optical channel ability 2</u>	45.2.1.133c
1.803	<u>Tx optical channel ability 3</u>	45.2.1.133d
1.804 through 1.819	Reserved	
1.820	<u>Rx optical channel control</u>	45.2.1.133e
1.821	<u>Rx optical channel ability 1</u>	45.2.1.133f
1.822	<u>Rx optical channel ability 2</u>	45.2.1.133g
1.823	<u>Rx optical channel ability 3</u>	45.2.1.133h
1.824 through 1.899	Reserved	
...		
1.2103 through 1.2293 1.2199	Reserved	
1.2200	<u>IFEC control</u>	45.2.1.186aa
1.2201	<u>IFEC status</u>	45.2.1.186ab
1.2202, 1.2203	<u>IFEC corrected codewords counter</u>	45.2.1.186ac
1.2204, 1.2205	<u>IFEC uncorrected codewords counter</u>	45.2.1.186ad
1.2206	<u>IFEC lane mapping</u>	45.2.1.186ae
1.2207 through 1.2209	Reserved	
1.2210 through 1.2217	<u>IFEC symbol error counter, lane 0 to 3</u>	45.2.1.186af, 45.2.1.186ag
1.2218 through 1.2245	Reserved	
1.2246	<u>SC-FEC alignment status 1</u>	45.2.1.186ah
1.2247	<u>SC-FEC alignment status 2</u>	45.2.1.186ai
1.2248, 1.2249	Reserved	

Table 45–3—PMA/PMD registers (continued)

Register address	Register name	Subclause
<u>1.2250 through 1.2269</u>	<u>SC-FEC lane mapping, lane 0 through 19</u>	<u>45.2.1.186aj</u> , <u>45.2.1.186ak</u>
<u>1.2270 through 1.2275</u>	<u>Reserved</u>	
<u>1.2276, 1.2277</u>	<u>SC-FEC corrected codewords counter</u>	<u>45.2.1.186al</u>
<u>1.2278, 1.2279</u>	<u>SC-FEC uncorrected codewords counter</u>	<u>45.2.1.186am</u>
<u>1.2280 through 1.2283</u>	<u>SC-FEC total bits</u>	<u>45.2.1.186an</u>
<u>1.2284 through 1.2287</u>	<u>SC-FEC corrected bits</u>	<u>45.2.1.186ao</u>
<u>1.2288 through 1.2293</u>	<u>Reserved</u>	
...		

45.2.1.6 PMA/PMD control 2 register (Register 1.7)

Change the indicated reserved row of Table 45–7 (as inserted by IEEE Std 802.3cu-2021) as follows (unchanged rows not shown):

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

Bit(s)	Name	Description	R/W ^a
...			
1.7.6:0	PMA/PMD type selection	6 5 4 3 2 1 0 ... 1 0 0 1 1 1 x = reserved 1 0 0 1 1 1 1 = reserved 1 0 0 1 1 1 0 = 100GBASE-ZR PMA/PMD ...	R/W

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

45.2.1.7 PMA/PMD status 2 register (Register 1.8)

45.2.1.7.4 Transmit fault (1.8.11)

Insert a row for 100GBASE-ZR in Table 45–9 after the row for “100GBASE-LR4, 100GBASE-ER4” as follows (unchanged rows not shown):

Table 45–9—Transmit fault description location

PMA/PMD	Description location
...	
100GBASE-ZR	154.5.8
...	

45.2.1.7.5 Receive fault (1.8.10)

Insert a row for 100GBASE-ZR in Table 45–10 after the row for “100GBASE-LR4, 100GBASE-ER4” as follows (unchanged rows not shown):

Table 45–10—Receive fault description location

PMA/PMD	Description location
...	
100GBASE-ZR	154.5.9
...	

45.2.1.8 PMD transmit disable register (Register 1.9)

Insert a row for 100GBASE-ZR in Table 45–12 after the row for “100GBASE-LR4 and 100GBASE-ER4” as follows (unchanged rows not shown):

Table 45–12—Transmit disable description location

PMA/PMD	Description location
...	
100GBASE-ZR	154.5.6
...	

45.2.1.21b 40G/100G PMA/PMD extended ability 2 register (Register 1.26)

Change the row for bit 1.26.6 in Table 45–24b (as modified by IEEE Std 802.3cu-2021) as follows (unchanged rows not shown):

Table 45–24b—40G/100G PMA/PMD extended ability 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
...			
1.26.6	Reserved 100GBASE-ZR ability	Value always 0 1 = PMA/PMD is able to perform 100GBASE-ZR 0 = PMA/PMD is not able to perform 100GBASE-ZR	RO
...			

^aRO = Read only

Insert 45.2.1.21b.3aa before 45.2.1.21b.3a (as inserted by IEEE Std 802.3cu-2021) as follows:

45.2.1.21b.3aa 100GBASE-ZR ability (1.26.6)

When read as a one, bit 1.26.6 indicates that the PMA/PMD is able to operate as a 100GBASE-ZR PMA/PMD type. When read as a zero, bit 1.26.6 indicates that the PMA/PMD is not able to operate as a 100GBASE-ZR PMA/PMD type.

Insert 45.2.1.133a through 45.2.1.133h after 45.2.1.133 as follows:

45.2.1.133a Tx optical channel control register (Register 1.800)

The assignment of bits in the Tx optical channel control register is shown in Table 45–102k.

Table 45–102k—Tx optical channel control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.800.15:6	Reserved	Value always 0	RO
1.800.5:0	Tx optical channel index	Integer value of the Tx optical channel index	R/W

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

45.2.1.133a.1 Tx optical channel index (1.800.5:0)

Bits 1.800.5:0 set the value of the Tx optical channel index (which directly relates to the optical channel and transmitter center frequency) with bit 1.800.0 being the LSB and bit 1.800.5 being the MSB. The channel index number indicates the corresponding optical frequency. For 100GBASE-ZR the specific optical frequency corresponding to each channel index number is listed in Table 154–5.

45.2.1.133b Tx optical channel ability 1 register (Register 1.801)

The assignment of bits in the Tx optical channel ability 1 register is shown in Table 45–102l.

Table 45–102l—Tx optical channel ability 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.801.15	Tx index ability 15	1 = PMD is able to operate with Tx optical channel index = 15 0 = PMD is not able to operate with Tx optical channel index = 15	RO
1.801.14	Tx index ability 14	1 = PMD is able to operate with Tx optical channel index = 14 0 = PMD is not able to operate with Tx optical channel index = 14	RO
1.801.13	Tx index ability 13	1 = PMD is able to operate with Tx optical channel index = 13 0 = PMD is not able to operate with Tx optical channel index = 13	RO
1.801.12	Tx index ability 12	1 = PMD is able to operate with Tx optical channel index = 12 0 = PMD is not able to operate with Tx optical channel index = 12	RO
1.801.11	Tx index ability 11	1 = PMD is able to operate with Tx optical channel index = 11 0 = PMD is not able to operate with Tx optical channel index = 11	RO
1.801.10	Tx index ability 10	1 = PMD is able to operate with Tx optical channel index = 10 0 = PMD is not able to operate with Tx optical channel index = 10	RO

Table 45–102l—Tx optical channel ability 1 register bit definitions (continued)

Bit(s)	Name	Description	R/W ^a
1.801.9	Tx index ability 9	1 = PMD is able to operate with Tx optical channel index = 9 0 = PMD is not able to operate with Tx optical channel index = 9	RO
1.801.8	Tx index ability 8	1 = PMD is able to operate with Tx optical channel index = 8 0 = PMD is not able to operate with Tx optical channel index = 8	RO
1.801.7	Tx index ability 7	1 = PMD is able to operate with Tx optical channel index = 7 0 = PMD is not able to operate with Tx optical channel index = 7	RO
1.801.6	Tx index ability 6	1 = PMD is able to operate with Tx optical channel index = 6 0 = PMD is not able to operate with Tx optical channel index = 6	RO
1.801.5	Tx index ability 5	1 = PMD is able to operate with Tx optical channel index = 5 0 = PMD is not able to operate with Tx optical channel index = 5	RO
1.801.4	Tx index ability 4	1 = PMD is able to operate with Tx optical channel index = 4 0 = PMD is not able to operate with Tx optical channel index = 4	RO
1.801.3	Tx index ability 3	1 = PMD is able to operate with Tx optical channel index = 3 0 = PMD is not able to operate with Tx optical channel index = 3	RO
1.801.2	Tx index ability 2	1 = PMD is able to operate with Tx optical channel index = 2 0 = PMD is not able to operate with Tx optical channel index = 2	RO
1.801.1	Tx index ability 1	1 = PMD is able to operate with Tx optical channel index = 1 0 = PMD is not able to operate with Tx optical channel index = 1	RO
1.801.0	Tx index ability 0	1 = PMD is able to operate with Tx optical channel index = 0 0 = PMD is not able to operate with Tx optical channel index = 0	RO

^aRO = Read only

45.2.1.133b.1 Tx index ability 0 through 15 (1.801.0 through 1.801.15)

When read as a one, bit 1.801.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with a Tx optical channel index of 0. When read as a zero, bit 1.801.0 indicates that the PMD is not able to operate with a Tx optical channel index of 0. Bits 1.801.1 through 1.801.15 indicate the equivalent for index values 1 through 15, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

45.2.1.133c Tx optical channel ability 2 register (Register 1.802)

The assignment of bits in the Tx optical channel ability 2 register is shown in Table 45–102m.

Table 45–102m—Tx optical channel ability 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.802.15	Tx index ability 31	1 = PMD is able to operate with Tx optical channel index = 31 0 = PMD is not able to operate with Tx optical channel index = 31	RO
1.802.14	Tx index ability 30	1 = PMD is able to operate with Tx optical channel index = 30 0 = PMD is not able to operate with Tx optical channel index = 30	RO
1.802.13	Tx index ability 29	1 = PMD is able to operate with Tx optical channel index = 29 0 = PMD is not able to operate with Tx optical channel index = 29	RO

Table 45–102m—Tx optical channel ability 2 register bit definitions (continued)

Bit(s)	Name	Description	R/W ^a
1.802.12	Tx index ability 28	1 = PMD is able to operate with Tx optical channel index = 28 0 = PMD is not able to operate with Tx optical channel index = 28	RO
1.802.11	Tx index ability 27	1 = PMD is able to operate with Tx optical channel index = 27 0 = PMD is not able to operate with Tx optical channel index = 27	RO
1.802.10	Tx index ability 26	1 = PMD is able to operate with Tx optical channel index = 26 0 = PMD is not able to operate with Tx optical channel index = 26	RO
1.802.9	Tx index ability 25	1 = PMD is able to operate with Tx optical channel index = 25 0 = PMD is not able to operate with Tx optical channel index = 25	RO
1.802.8	Tx index ability 24	1 = PMD is able to operate with Tx optical channel index = 24 0 = PMD is not able to operate with Tx optical channel index = 24	RO
1.802.7	Tx index ability 23	1 = PMD is able to operate with Tx optical channel index = 23 0 = PMD is not able to operate with Tx optical channel index = 23	RO
1.802.6	Tx index ability 22	1 = PMD is able to operate with Tx optical channel index = 22 0 = PMD is not able to operate with Tx optical channel index = 22	RO
1.802.5	Tx index ability 21	1 = PMD is able to operate with Tx optical channel index = 21 0 = PMD is not able to operate with Tx optical channel index = 21	RO
1.802.4	Tx index ability 20	1 = PMD is able to operate with Tx optical channel index = 20 0 = PMD is not able to operate with Tx optical channel index = 20	RO
1.802.3	Tx index ability 19	1 = PMD is able to operate with Tx optical channel index = 19 0 = PMD is not able to operate with Tx optical channel index = 19	RO
1.802.2	Tx index ability 18	1 = PMD is able to operate with Tx optical channel index = 18 0 = PMD is not able to operate with Tx optical channel index = 18	RO
1.802.1	Tx index ability 17	1 = PMD is able to operate with Tx optical channel index = 17 0 = PMD is not able to operate with Tx optical channel index = 17	RO
1.802.0	Tx index ability 16	1 = PMD is able to operate with Tx optical channel index = 16 0 = PMD is not able to operate with Tx optical channel index = 16	RO

^aRO = Read only**45.2.1.133c.1 Tx index ability 16 through 31 (1.802.0 through 1.802.15)**

When read as a one, bit 1.802.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with a Tx optical channel index of 16. When read as a zero, bit 1.802.0 indicates that the PMD is not able to operate with a Tx optical channel index of 16. Bits 1.802.1 through 1.802.15 indicate the equivalent for index values 17 through 31, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

45.2.1.133d Tx optical channel ability 3 register (Register 1.803)

The assignment of bits in the Tx optical channel ability 3 register is shown in Table 45–102n.

Table 45–102n—Tx optical channel ability 3 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.803.15	Tx index ability 47	1 = PMD is able to operate with Tx optical channel index = 47 0 = PMD is not able to operate with Tx optical channel index = 47	RO
1.803.14	Tx index ability 46	1 = PMD is able to operate with Tx optical channel index = 46 0 = PMD is not able to operate with Tx optical channel index = 46	RO
1.803.13	Tx index ability 45	1 = PMD is able to operate with Tx optical channel index = 45 0 = PMD is not able to operate with Tx optical channel index = 45	RO
1.803.12	Tx index ability 44	1 = PMD is able to operate with Tx optical channel index = 44 0 = PMD is not able to operate with Tx optical channel index = 44	RO
1.803.11	Tx index ability 43	1 = PMD is able to operate with Tx optical channel index = 43 0 = PMD is not able to operate with Tx optical channel index = 43	RO
1.803.10	Tx index ability 42	1 = PMD is able to operate with Tx optical channel index = 42 0 = PMD is not able to operate with Tx optical channel index = 42	RO
1.803.9	Tx index ability 41	1 = PMD is able to operate with Tx optical channel index = 41 0 = PMD is not able to operate with Tx optical channel index = 41	RO
1.803.8	Tx index ability 40	1 = PMD is able to operate with Tx optical channel index = 40 0 = PMD is not able to operate with Tx optical channel index = 40	RO
1.803.7	Tx index ability 39	1 = PMD is able to operate with Tx optical channel index = 39 0 = PMD is not able to operate with Tx optical channel index = 39	RO
1.803.6	Tx index ability 38	1 = PMD is able to operate with Tx optical channel index = 38 0 = PMD is not able to operate with Tx optical channel index = 38	RO
1.803.5	Tx index ability 37	1 = PMD is able to operate with Tx optical channel index = 37 0 = PMD is not able to operate with Tx optical channel index = 37	RO
1.803.4	Tx index ability 36	1 = PMD is able to operate with Tx optical channel index = 36 0 = PMD is not able to operate with Tx optical channel index = 36	RO
1.803.3	Tx index ability 35	1 = PMD is able to operate with Tx optical channel index = 35 0 = PMD is not able to operate with Tx optical channel index = 35	RO
1.803.2	Tx index ability 34	1 = PMD is able to operate with Tx optical channel index = 34 0 = PMD is not able to operate with Tx optical channel index = 34	RO
1.803.1	Tx index ability 33	1 = PMD is able to operate with Tx optical channel index = 33 0 = PMD is not able to operate with Tx optical channel index = 33	RO
1.803.0	Tx index ability 32	1 = PMD is able to operate with Tx optical channel index = 32 0 = PMD is not able to operate with Tx optical channel index = 32	RO

^aRO = Read only

45.2.1.133d.1 Tx index ability 32 through 47 (1.803.0 through 1.803.15)

When read as a one, bit 1.803.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with a Tx optical channel index of 32. When read as a zero, bit 1.803.0 indicates that the PMD is not able to operate with a Tx optical channel index of 32. Bits 1.803.1 through 1.803.15 indicate the equivalent for index values

33 through 47, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

45.2.1.133e Rx optical channel control register (Register 1.820)

The assignment of bits in the Rx optical channel control register is shown in Table 45–102o.

Table 45–102o—Rx optical channel control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.820.15	Tx Rx different optical channel ability	1 = PMD is able to operate with different Tx and Rx optical channels 0 = PMD is not able to operate with different Tx and Rx optical channels	RO
1.820.14:6	Reserved	Value always 0	RO
1.820.5:0	Rx optical channel index	Integer value of the Rx optical channel index	R/W

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

45.2.1.133e.1 Tx Rx different optical channel ability (1.820.15)

When read as a one, bit 1.820.15 indicates that the PMD is able to operate with an Rx optical channel index that is different from the Tx optical channel index. When read as a zero, bit 1.820.15 indicates that the PMD is not able to operate with an Rx optical channel index that is different from the Tx optical channel index.

45.2.1.133e.2 Rx optical channel index (1.820.5:0)

If the PMD is able to operate with an Rx optical channel index number that is different from the Tx optical channel index number (bit 1.820.15 is one), bits 1.820.5:0 set the value of the Rx optical channel index number (and hence the receiver optical frequency) with bit 1.820.0 being the LSB and bit 1.820.5 being the MSB. The channel index number indicates the corresponding optical frequency. For 100GBASE-ZR the specific optical frequency corresponding to each channel index number is listed in Table 154–5. If the PMD is not able to operate with an Rx optical channel index that is different from the Tx optical channel index (bit 1.820.15 is zero), bits 1.820.5:0 have no effect and the value of the Rx optical channel index is the same as the value of the Tx optical channel index as set by bits 1.800.5:0.

45.2.1.133f Rx optical channel ability 1 register (Register 1.821)

The assignment of bits in the Rx optical channel ability 1 register is shown in Table 45–102p.

Table 45–102p—Rx optical channel ability 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.821.15	Rx index ability 15	1 = PMD is able to operate with Rx optical channel index = 15 0 = PMD is not able to operate with Rx optical channel index = 15	RO
1.821.14	Rx index ability 14	1 = PMD is able to operate with Rx optical channel index = 14 0 = PMD is not able to operate with Rx optical channel index = 14	RO
1.821.13	Rx index ability 13	1 = PMD is able to operate with Rx optical channel index = 13 0 = PMD is not able to operate with Rx optical channel index = 13	RO
1.821.12	Rx index ability 12	1 = PMD is able to operate with Rx optical channel index = 12 0 = PMD is not able to operate with Rx optical channel index = 12	RO
1.821.11	Rx index ability 11	1 = PMD is able to operate with Rx optical channel index = 11 0 = PMD is not able to operate with Rx optical channel index = 11	RO
1.821.10	Rx index ability 10	1 = PMD is able to operate with Rx optical channel index = 10 0 = PMD is not able to operate with Rx optical channel index = 10	RO
1.821.9	Rx index ability 9	1 = PMD is able to operate with Rx optical channel index = 9 0 = PMD is not able to operate with Rx optical channel index = 9	RO
1.821.8	Rx index ability 8	1 = PMD is able to operate with Rx optical channel index = 8 0 = PMD is not able to operate with Rx optical channel index = 8	RO
1.821.7	Rx index ability 7	1 = PMD is able to operate with Rx optical channel index = 7 0 = PMD is not able to operate with Rx optical channel index = 7	RO
1.821.6	Rx index ability 6	1 = PMD is able to operate with Rx optical channel index = 6 0 = PMD is not able to operate with Rx optical channel index = 6	RO
1.821.5	Rx index ability 5	1 = PMD is able to operate with Rx optical channel index = 5 0 = PMD is not able to operate with Rx optical channel index = 5	RO
1.821.4	Rx index ability 4	1 = PMD is able to operate with Rx optical channel index = 4 0 = PMD is not able to operate with Rx optical channel index = 4	RO
1.821.3	Rx index ability 3	1 = PMD is able to operate with Rx optical channel index = 3 0 = PMD is not able to operate with Rx optical channel index = 3	RO
1.821.2	Rx index ability 2	1 = PMD is able to operate with Rx optical channel index = 2 0 = PMD is not able to operate with Rx optical channel index = 2	RO
1.821.1	Rx index ability 1	1 = PMD is able to operate with Rx optical channel index = 1 0 = PMD is not able to operate with Rx optical channel index = 1	RO
1.821.0	Rx index ability 0	1 = PMD is able to operate with Rx optical channel index = 0 0 = PMD is not able to operate with Rx optical channel index = 0	RO

^aRO = Read only

45.2.1.133f.1 Rx index ability 0 through 15 (1.821.0 through 1.821.15)

When read as a one, bit 1.821.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with an Rx optical channel index of 0. When read as a zero, bit 1.821.0 indicates that the PMD is not able to operate with an Rx optical channel index of 0. Bits 1.821.1 through 1.821.15 indicate the equivalent for index values

1 through 15, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

45.2.1.133g Rx optical channel ability 2 register (Register 1.822)

The assignment of bits in the Rx optical channel ability 2 register is shown in Table 45–102q.

Table 45–102q—Rx optical channel ability 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.822.15	Rx index ability 31	1 = PMD is able to operate with Rx optical channel index = 31 0 = PMD is not able to operate with Rx optical channel index = 31	RO
1.822.14	Rx index ability 30	1 = PMD is able to operate with Rx optical channel index = 30 0 = PMD is not able to operate with Rx optical channel index = 30	RO
1.822.13	Rx index ability 29	1 = PMD is able to operate with Rx optical channel index = 29 0 = PMD is not able to operate with Rx optical channel index = 29	RO
1.822.12	Rx index ability 28	1 = PMD is able to operate with Rx optical channel index = 28 0 = PMD is not able to operate with Rx optical channel index = 28	RO
1.822.11	Rx index ability 27	1 = PMD is able to operate with Rx optical channel index = 27 0 = PMD is not able to operate with Rx optical channel index = 27	RO
1.822.10	Rx index ability 26	1 = PMD is able to operate with Rx optical channel index = 26 0 = PMD is not able to operate with Rx optical channel index = 26	RO
1.822.9	Rx index ability 25	1 = PMD is able to operate with Rx optical channel index = 25 0 = PMD is not able to operate with Rx optical channel index = 25	RO
1.822.8	Rx index ability 24	1 = PMD is able to operate with Rx optical channel index = 24 0 = PMD is not able to operate with Rx optical channel index = 24	RO
1.822.7	Rx index ability 23	1 = PMD is able to operate with Rx optical channel index = 23 0 = PMD is not able to operate with Rx optical channel index = 23	RO
1.822.6	Rx index ability 22	1 = PMD is able to operate with Rx optical channel index = 22 0 = PMD is not able to operate with Rx optical channel index = 22	RO
1.822.5	Rx index ability 21	1 = PMD is able to operate with Rx optical channel index = 21 0 = PMD is not able to operate with Rx optical channel index = 21	RO
1.822.4	Rx index ability 20	1 = PMD is able to operate with Rx optical channel index = 20 0 = PMD is not able to operate with Rx optical channel index = 20	RO
1.822.3	Rx index ability 19	1 = PMD is able to operate with Rx optical channel index = 19 0 = PMD is not able to operate with Rx optical channel index = 19	RO
1.822.2	Rx index ability 18	1 = PMD is able to operate with Rx optical channel index = 18 0 = PMD is not able to operate with Rx optical channel index = 18	RO
1.822.1	Rx index ability 17	1 = PMD is able to operate with Rx optical channel index = 17 0 = PMD is not able to operate with Rx optical channel index = 17	RO
1.822.0	Rx index ability 16	1 = PMD is able to operate with Rx optical channel index = 16 0 = PMD is not able to operate with Rx optical channel index = 16	RO

^aRO = Read only

45.2.1.133g.1 Rx index ability 16 through 31 (1.822.0 through 1.822.15)

When read as a one, bit 1.822.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with an Rx optical channel index of 16. When read as a zero, bit 1.822.0 indicates that the PMD is not able to operate with an Rx optical channel index of 16. Bits 1.822.1 through 1.822.15 indicate the equivalent for index values 17 through 31, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

45.2.1.133h Rx optical channel ability 3 register (Register 1.823)

The assignment of bits in the Rx optical channel ability 3 register is shown in Table 45–102r.

Table 45–102r—Rx optical channel ability 3 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.823.15	Rx index ability 47	1 = PMD is able to operate with Rx optical channel index = 47 0 = PMD is not able to operate with Rx optical channel index = 47	RO
1.823.14	Rx index ability 46	1 = PMD is able to operate with Rx optical channel index = 46 0 = PMD is not able to operate with Rx optical channel index = 46	RO
1.823.13	Rx index ability 45	1 = PMD is able to operate with Rx optical channel index = 45 0 = PMD is not able to operate with Rx optical channel index = 45	RO
1.823.12	Rx index ability 44	1 = PMD is able to operate with Rx optical channel index = 44 0 = PMD is not able to operate with Rx optical channel index = 44	RO
1.823.11	Rx index ability 43	1 = PMD is able to operate with Rx optical channel index = 43 0 = PMD is not able to operate with Rx optical channel index = 43	RO
1.823.10	Rx index ability 42	1 = PMD is able to operate with Rx optical channel index = 42 0 = PMD is not able to operate with Rx optical channel index = 42	RO
1.823.9	Rx index ability 41	1 = PMD is able to operate with Rx optical channel index = 41 0 = PMD is not able to operate with Rx optical channel index = 41	RO
1.823.8	Rx index ability 40	1 = PMD is able to operate with Rx optical channel index = 40 0 = PMD is not able to operate with Rx optical channel index = 40	RO
1.823.7	Rx index ability 39	1 = PMD is able to operate with Rx optical channel index = 39 0 = PMD is not able to operate with Rx optical channel index = 39	RO
1.823.6	Rx index ability 38	1 = PMD is able to operate with Rx optical channel index = 38 0 = PMD is not able to operate with Rx optical channel index = 38	RO
1.823.5	Rx index ability 37	1 = PMD is able to operate with Rx optical channel index = 37 0 = PMD is not able to operate with Rx optical channel index = 37	RO
1.823.4	Rx index ability 36	1 = PMD is able to operate with Rx optical channel index = 36 0 = PMD is not able to operate with Rx optical channel index = 36	RO
1.823.3	Rx index ability 35	1 = PMD is able to operate with Rx optical channel index = 35 0 = PMD is not able to operate with Rx optical channel index = 35	RO
1.823.2	Rx index ability 34	1 = PMD is able to operate with Rx optical channel index = 34 0 = PMD is not able to operate with Rx optical channel index = 34	RO
1.823.1	Rx index ability 33	1 = PMD is able to operate with Rx optical channel index = 33 0 = PMD is not able to operate with Rx optical channel index = 33	RO
1.823.0	Rx index ability 32	1 = PMD is able to operate with Rx optical channel index = 32 0 = PMD is not able to operate with Rx optical channel index = 32	RO

^aRO = Read only

45.2.1.133h.1 Rx index ability 32 through 47 (1.823.0 through 1.823.15)

When read as a one, bit 1.823.0 indicates that the PMD selected by bits 1.7.6:0 is able to operate with an Rx optical channel index of 32. When read as a zero, bit 1.823.0 indicates that the PMD is not able to operate with an Rx optical channel index of 32. Bits 1.823.1 through 1.823.15 indicate the equivalent for index values 33 through 47, respectively. The optical frequencies that correspond to these index values are given in the appropriate PMD clause. For 100GBASE-ZR see Table 154–5.

Insert 45.2.1.186aa through 45.2.1.186ao before 45.2.1.186a (as inserted by IEEE Std 802.3cg-2019) as follows:

45.2.1.186aa IFEC control register (Register 1.2200)

The assignment of bits in the IFEC control register is shown in Table 45–150aa.

Table 45–150aa—IFEC control register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2200.15:2	Reserved	Value always 0	RO
1.2200.1	IFEC bypass indication enable	1 = IFEC decoder does not indicate errors 0 = IFEC decoder indicates errors	R/W
1.2200.0	IFEC bypass correction enable	1 = IFEC decoder performs error detection without error correction 0 = IFEC decoder performs error detection and error correction	R/W

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

45.2.1.186aa.1 IFEC bypass indication enable (1.2200.1)

This bit enables the Inverse RS-FEC decoder to bypass error indication to the far end (PCS) through the sync bits for the BASE-R PHY in the remote device. When set to one, this bit enables bypass of the error indication function. When set to zero, errors are indicated to the remote PCS through the sync bits. Writes to bit 1.2200.1 are ignored and reads return a zero if the Inverse RS-FEC does not have the ability to bypass decoding error indications to the remote PCS layer (see 152.5.2.3).

45.2.1.186aa.2 IFEC bypass correction enable (1.2200.0)

When bit 1.2200.0 is set to one the Inverse RS-FEC decoder performs error detection without error correction (see 152.5.2.3). When this bit is set to zero, the decoder also performs error correction. Writes to this bit are ignored and reads return a zero if the Inverse RS-FEC does not have the ability to bypass error correction.

45.2.1.186ab IFEC status register (Register 1.2201)

The assignment of bits in the IFEC status register is shown in Table 45–150ab.

Table 45–150ab—IFEC status register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2201.15	PCS align status	1 = IFEC encoder has locked and aligned all PCS lanes 0 = IFEC encoder has not locked and aligned all PCS lanes	RO
1.2201.14	IFEC align status	1 = IFEC receive lanes locked and aligned 0 = IFEC receive lanes not locked and aligned	RO
1.2201.13:12	Reserved	Value always 0	RO
1.2201.11	IFEC AM lock 3	1 = IFEC lane 3 locked and aligned 0 = IFEC lane 3 not locked and aligned	RO
1.2201.10	IFEC AM lock 2	1 = IFEC lane 2 locked and aligned 0 = IFEC lane 2 not locked and aligned	RO
1.2201.9	IFEC AM lock 1	1 = IFEC lane 1 locked and aligned 0 = IFEC lane 1 not locked and aligned	RO
1.2201.8	IFEC AM lock 0	1 = IFEC lane 0 locked and aligned 0 = IFEC lane 0 not locked and aligned	RO
1.2201.7:3	Reserved	Value always 0	RO
1.2201.2	IFEC high SER	1 = IFEC errors have exceeded threshold 0 = IFEC errors have not exceeded threshold	RO/LH
1.2201.1	IFEC bypass indication ability	1 = IFEC decoder has the ability to bypass error indication 0 = IFEC decoder does not have the ability to bypass error indication	RO
1.2201.0	IFEC bypass correction ability	1 = IFEC decoder has the ability to bypass error correction 0 = IFEC decoder does not have the ability to bypass error correction	RO

^aRO = Read only, LH = Latching high

45.2.1.186ab.1 PCS align status (1.2201.15)

Bit 1.2201.15 indicates the PCS alignment status of the Inverse RS-FEC. For the Inverse RS-FEC described in Clause 152, PCS alignment is defined as block lock, alignment marker lock, and deskew of all 20 receive PCS lanes. When read as a zero, bit 1.2201.15 indicates that the Inverse RS-FEC has not obtained PCS alignment. When read as a one, bit 1.2201.15 indicates that the Inverse RS-FEC has obtained PCS alignment.

45.2.1.186ab.2 IFEC align status (1.2201.14)

Bit 1.2201.14 indicates the alignment status of the Inverse RS-FEC. For the Inverse RS-FEC described in Clause 152, IFEC alignment is defined as alignment marker lock and deskew of all four lanes on the IFEC service interface. When read as a zero, bit 1.2201.14 indicates that the Inverse RS-FEC has not obtained IFEC alignment. When read as a one, bit 1.2201.14 indicates that the Inverse RS-FEC has obtained IFEC alignment.

45.2.1.186ab.3 IFEC AM lock 3 (1.2201.11)

When read as a one, bit 1.2201.11 indicates that the Inverse RS-FEC described in Clause 152 has locked and aligned lane 3 of the IFEC service interface. When read as a zero, bit 1.2201.11 indicates that the Inverse

RS-FEC has not locked and aligned lane 3 of the IFEC service interface. This bit reflects the state of `amps_lock[3]` (see 152.5.2.1).

45.2.1.186ab.4 IFEC AM lock 2 (1.2201.10)

When read as a one, bit 1.2201.10 indicates that the Inverse RS-FEC described in Clause 152 has locked and aligned lane 2 of the IFEC service interface. When read as a zero, bit 1.2201.10 indicates that the Inverse RS-FEC has not locked and aligned lane 2 of the IFEC service interface. This bit reflects the state of `amps_lock[2]` (see 152.5.2.1).

45.2.1.186ab.5 IFEC AM lock 1 (1.2201.9)

When read as a one, bit 1.2201.9 indicates that the Inverse RS-FEC described in Clause 152 has locked and aligned lane 1 of the IFEC service interface. When read as a zero, bit 1.2201.9 indicates that the Inverse RS-FEC has not locked and aligned lane 1 of the IFEC service interface. This bit reflects the state of `amps_lock[1]` (see 152.5.2.1).

45.2.1.186ab.6 IFEC AM lock 0 (1.2201.8)

When read as a one, bit 1.2201.8 indicates that the Inverse RS-FEC described in Clause 152 has locked and aligned lane 0 of the IFEC service interface. When read as a zero, bit 1.2201.8 indicates that the Inverse RS-FEC has not locked and aligned lane 0 of the IFEC service interface. This bit reflects the state of `amps_lock[0]` (see 152.5.2.1).

45.2.1.186ab.7 IFEC high SER (1.2201.2)

When bit 1.2200.1 (IFEC bypass indication enable) is set to one, bit 1.2201.2 is set to one if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 152.5.2.3) and is set to zero otherwise. This bit is set to zero if bit 1.2200.1 (IFEC bypass indication enable) is set to zero. This bit shall be implemented with latching high behavior.

45.2.1.186ab.8 IFEC bypass indication ability (1.2201.1)

The Reed-Solomon decoder may have the option to perform error detection without error indication (see 152.5.2.3) to reduce the delay contributed by the Inverse RS-FEC sublayer. This bit is set to one to indicate that the decoder has this ability to bypass the error indication function. The bit is set to zero if this ability is not supported.

45.2.1.186ab.9 IFEC bypass correction ability (1.2201.0)

The Reed-Solomon decoder may have the option to perform error detection without error correction (see 152.5.2.3) to reduce the delay contributed by the Inverse RS-FEC sublayer. This bit is set to one to indicate that the decoder has this ability to bypass error correction. The bit is set to zero if this ability is not supported.

45.2.1.186ac IFEC corrected codewords counter (Register 1.2202, 1.2203)

The assignment of bits in the IFEC corrected codewords counter is shown in Table 45–150ac. See 152.6.8 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2202 and 1.2203 are used to read the value of a 32-bit counter. When registers 1.2202 and 1.2203 are used to read the 32-bit counter value, register 1.2202 is read first, the value of the register 1.2203

is latched when (and only when) register 1.2202 is read, and reads of register 1.2203 return the latched value rather than the current value of the counter.

Table 45–150ac—IFEC corrected codewords counter bit definitions

Bit(s)	Name	Description	R/W ^a
1.2202.15:0	IFEC corrected codewords lower	IFEC_corrected_cw_counter[15:0]	RO, NR
1.2203.15:0	IFEC corrected codewords upper	IFEC_corrected_cw_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

45.2.1.186ad IFEC uncorrected codewords counter (Register 1.2204, 1.2205)

The assignment of bits in the IFEC uncorrected codewords counter is shown in Table 45–150ad. See 152.6.9 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2204 and 1.2205 are used to read the value of a 32-bit counter. When registers 1.2204 and 1.2205 are used to read the 32-bit counter value, register 1.2204 is read first, the value of the register 1.2205 is latched when (and only when) register 1.2204 is read, and reads of register 1.2205 return the latched value rather than the current value of the counter.

Table 45–150ad—IFEC uncorrected codewords counter bit definitions

Bit(s)	Name	Description	R/W ^a
1.2204.15:0	IFEC uncorrected codewords lower	IFEC_uncorrected_cw_counter[15:0]	RO, NR
1.2205.15:0	IFEC uncorrected codewords upper	IFEC_uncorrected_cw_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

45.2.1.186ae IFEC lane mapping register (Register 1.2206)

The assignment of bits in the IFEC lane mapping register is shown in Table 45–150ae. When the RS-FEC detects and locks the RS-FEC for IFEC service interface lane 0, the detected RS-FEC lane number is recorded in bits 1:0 in this register. Similarly, the detected RS-FEC lane numbers for IFEC service lanes 1, 2, and 3 are recorded in bits 3:2, 5:4, and 7:6, respectively. The contents of the RS-FEC lane mapping register bits 7:0 are valid when Inverse RS-FEC align status (1.2201.14) is set to one and are invalid otherwise.

Table 45–150ae—IFEC lane mapping register

Bit(s)	Name	Description	R/W ^a
1.2206.15:8	Reserved	Value always 0	RO
1.2206.7:6	RS-FEC lane 3 mapping	RS-FEC lane mapped to IFEC lane 3	RO
1.2206.5:4	RS-FEC lane 2 mapping	RS-FEC lane mapped to IFEC lane 2	RO

Table 45–150ae—IFEC lane mapping register (continued)

Bit(s)	Name	Description	R/W ^a
1.2206.3:2	RS-FEC lane 1 mapping	RS-FEC lane mapped to IFEC lane 1	RO
1.2206.1:0	RS-FEC lane 0 mapping	RS-FEC lane mapped to IFEC lane 0	RO

^aRO = Read only**45.2.1.186af IFEC symbol error counter, lane 0 (Register 1.2210, 1.2211)**

The assignment of bits in the IFEC symbol error counter, lane 0 is shown in Table 45–150af. Symbol errors detected in FEC lane 0 are counted and shown in bits 1.2210.15:0 and 1.2211.15:0. See 152.6.11 for a definition of this counter. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2210 and 1.2211 are used to read the value of a 32-bit counter. When registers 1.2210 and 1.2211 are used to read the 32-bit counter value, register 1.2210 is read first, the value of the register 1.2211 is latched when (and only when) register 1.2210 is read, and reads of register 1.2211 return the latched value rather than the current value of the counter.

Table 45–150af—IFEC symbol error counter, lane 0 bit definitions

Bit(s)	Name	Description	R/W ^a
1.2210.15:0	IFEC symbol errors, lane 0 lower	IFEC_symbol_error_counter_0[15:0]	RO, NR
1.2211.15:0	IFEC symbol errors, lane 0 upper	IFEC_symbol_error_counter_0[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over**45.2.1.186ag IFEC symbol error counter, lane 1 through 3 (Register 1.2212, 1.2213, 1.2214, 1.2215, 1.2216, 1.2217)**

The behavior of the IFEC symbol error counters, lane 1 through 3 is identical to that described for FEC lane 0 in 45.2.1.186af. Errors detected in each FEC lane are counted and shown in the corresponding register. FEC lane 1, lower 16 bits are shown in register 1.2212; FEC lane 1, upper 16 bits are shown in register 1.2213; FEC lane 2, lower 16 bits are shown in register 1.2214; through register 1.2217 for FEC lane 3, upper 16 bits.

45.2.1.186ah SC-FEC alignment status 1 register (Register 1.2246)

The assignment of bits in the staircase FEC (SC-FEC) alignment status 1 register is shown in Table 45–150ag.

45.2.1.186ah.1 SC-FEC align status (1.2246.12)

When read as a one, bit 1.2246.12 indicates that the SC-FEC has locked and aligned all receive lanes. When read as a zero, bit 1.2246.12 indicates that the SC-FEC has not locked and aligned all receive lanes.

Table 45–150ag—SC-FEC alignment status 1 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2246.15:13	Reserved	Value always 0	RO
1.2246.12	SC-FEC align status	1 = SC-FEC receive lanes locked and aligned 0 = SC-FEC receive lanes not locked and aligned	RO
1.2246.11:8	Reserved	Value always 0	RO
1.2246.7	SC-FEC FAS lock 7	1 = Lane 7 is locked 0 = Lane 7 is not locked	RO
1.2246.6	SC-FEC FAS lock 6	1 = Lane 6 is locked 0 = Lane 6 is not locked	RO
1.2246.5	SC-FEC FAS lock 5	1 = Lane 5 is locked 0 = Lane 5 is not locked	RO
1.2246.4	SC-FEC FAS lock 4	1 = Lane 4 is locked 0 = Lane 4 is not locked	RO
1.2246.3	SC-FEC FAS lock 3	1 = Lane 3 is locked 0 = Lane 3 is not locked	RO
1.2246.2	SC-FEC FAS lock 2	1 = Lane 2 is locked 0 = Lane 2 is not locked	RO
1.2246.1	SC-FEC FAS lock 1	1 = Lane 1 is locked 0 = Lane 1 is not locked	RO
1.2246.0	SC-FEC FAS lock 0	1 = Lane 0 is locked 0 = Lane 0 is not locked	RO

^aRO = Read only**45.2.1.186ah.2 SC-FEC FAS lock 7 (1.2246.7)**

When read as a one, bit 1.2246.7 indicates that the SC-FEC receiver has achieved frame alignment signal (FAS) lock for lane 7 of the PMA service interface. When read as a zero, bit 1.2246.7 indicates that the SC-FEC receiver has not achieved FAS lock for lane 7 of the PMA service interface. This bit reflects the state of fas_lock<7> (see 153.2.4.1.1).

45.2.1.186ah.3 SC-FEC FAS lock 6 (1.2246.6)

When read as a one, bit 1.2246.6 indicates that the SC-FEC receiver has achieved FAS lock for lane 6 of the PMA service interface. When read as a zero, bit 1.2246.6 indicates that the SC-FEC receiver has not achieved FAS lock for lane 6 of the PMA service interface. This bit reflects the state of fas_lock<6> (see 153.2.4.1.1).

45.2.1.186ah.4 SC-FEC FAS lock 5 (1.2246.5)

When read as a one, bit 1.2246.5 indicates that the SC-FEC receiver has achieved FAS lock for lane 5 of the PMA service interface. When read as a zero, bit 1.2246.5 indicates that the SC-FEC receiver has not

achieved FAS lock for lane 5 of the PMA service interface. This bit reflects the state of fas_lock<5> (see 153.2.4.1.1).

45.2.1.186ah.5 SC-FEC FAS lock 4 (1.2246.4)

When read as a one, bit 1.2246.4 indicates that the SC-FEC receiver has achieved FAS lock for lane 4 of the PMA service interface. When read as a zero, bit 1.2246.4 indicates that the SC-FEC receiver has not achieved FAS lock for lane 4 of the PMA service interface. This bit reflects the state of fas_lock<4> (see 153.2.4.1.1).

45.2.1.186ah.6 SC-FEC FAS lock 3 (1.2246.3)

When read as a one, bit 1.2246.3 indicates that the SC-FEC receiver has achieved FAS lock for lane 3 of the PMA service interface. When read as a zero, bit 1.2246.3 indicates that the SC-FEC receiver has not achieved FAS lock for lane 3 of the PMA service interface. This bit reflects the state of fas_lock<3> (see 153.2.4.1.1).

45.2.1.186ah.7 SC-FEC FAS lock 2 (1.2246.2)

When read as a one, bit 1.2246.2 indicates that the SC-FEC receiver has achieved FAS lock for lane 2 of the PMA service interface. When read as a zero, bit 1.2246.2 indicates that the SC-FEC receiver has not achieved FAS lock for lane 2 of the PMA service interface. This bit reflects the state of fas_lock<2> (see 153.2.4.1.1).

45.2.1.186ah.8 SC-FEC FAS lock 1 (1.2246.1)

When read as a one, bit 1.2246.1 indicates that the SC-FEC receiver has achieved FAS lock for lane 1 of the PMA service interface. When read as a zero, bit 1.2246.1 indicates that the SC-FEC receiver has not achieved FAS lock for lane 1 of the PMA service interface. This bit reflects the state of fas_lock<1> (see 153.2.4.1.1).

45.2.1.186ah.9 SC-FEC FAS lock 0 (1.2246.0)

When read as a one, bit 1.2246.0 indicates that the SC-FEC receiver has achieved FAS lock for lane 0 of the PMA service interface. When read as a zero, bit 1.2246.0 indicates that the SC-FEC receiver has not achieved FAS lock for lane 0 of the PMA service interface. This bit reflects the state of fas_lock<0> (see 153.2.4.1.1).

45.2.1.186ai SC-FEC alignment status 2 register (Register 1.2247)

The assignment of bits in the SC-FEC alignment status 2 register is shown in Table 45–150ah.

Table 45–150ah—SC-FEC alignment status 2 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2247.15:12	Reserved	Value always 0	RO
1.2247.11	SC-FEC FAS lock 19	1 = Lane 19 is locked 0 = Lane 19 is not locked	RO
1.2247.10	SC-FEC FAS lock 18	1 = Lane 18 is locked 0 = Lane 18 is not locked	RO

Table 45–150ah—SC-FEC alignment status 2 register bit definitions (continued)

Bit(s)	Name	Description	R/W ^a
1.2247.9	SC-FEC FAS lock 17	1 = Lane 17 is locked 0 = Lane 17 is not locked	RO
1.2247.8	SC-FEC FAS lock 16	1 = Lane 16 is locked 0 = Lane 16 is not locked	RO
1.2247.7	SC-FEC FAS lock 15	1 = Lane 15 is locked 0 = Lane 15 is not locked	RO
1.2247.6	SC-FEC FAS lock 14	1 = Lane 14 is locked 0 = Lane 14 is not locked	RO
1.2247.5	SC-FEC FAS lock 13	1 = Lane 13 is locked 0 = Lane 13 is not locked	RO
1.2247.4	SC-FEC FAS lock 12	1 = Lane 12 is locked 0 = Lane 12 is not locked	RO
1.2247.3	SC-FEC FAS lock 11	1 = Lane 11 is locked 0 = Lane 11 is not locked	RO
1.2247.2	SC-FEC FAS lock 10	1 = Lane 10 is locked 0 = Lane 10 is not locked	RO
1.2247.1	SC-FEC FAS lock 9	1 = Lane 9 is locked 0 = Lane 9 is not locked	RO
1.2247.0	SC-FEC FAS lock 8	1 = Lane 8 is locked 0 = Lane 8 is not locked	RO

^aRO = Read only

45.2.1.186ai.1 SC-FEC FAS lock 19 (1.2247.11)

When read as a one, bit 1.2247.11 indicates that the SC-FEC receiver has achieved FAS lock for lane 19 of the PMA service interface. When read as a zero, bit 1.2247.11 indicates that the SC-FEC receiver has not achieved FAS lock for lane 19 of the PMA service interface. This bit reflects the state of fas_lock<19> (see 153.2.4.1.1).

45.2.1.186ai.2 SC-FEC FAS lock 18 (1.2247.10)

When read as a one, bit 1.2247.10 indicates that the SC-FEC receiver has achieved FAS lock for lane 18 of the PMA service interface. When read as a zero, bit 1.2247.10 indicates that the SC-FEC receiver has not achieved FAS lock for lane 18 of the PMA service interface. This bit reflects the state of fas_lock<18> (see 153.2.4.1.1).

45.2.1.186ai.3 SC-FEC FAS lock 17 (1.2247.9)

When read as a one, bit 1.2247.9 indicates that the SC-FEC receiver has achieved FAS lock for lane 17 of the PMA service interface. When read as a zero, bit 1.2247.9 indicates that the SC-FEC receiver has not achieved FAS lock for lane 17 of the PMA service interface. This bit reflects the state of fas_lock<17> (see 153.2.4.1.1).

45.2.1.186ai.4 SC-FEC FAS lock 16 (1.2247.8)

When read as a one, bit 1.2247.8 indicates that the SC-FEC receiver has achieved FAS lock for lane 16 of the PMA service interface. When read as a zero, bit 1.2247.8 indicates that the SC-FEC receiver has not achieved FAS lock for lane 16 of the PMA service interface This bit reflects the state of fas_lock<16> (see 153.2.4.1.1).

45.2.1.186ai.5 SC-FEC FAS lock 15 (1.2247.7)

When read as a one, bit 1.2247.7 indicates that the SC-FEC receiver has achieved FAS lock for lane 15 of the PMA service interface. When read as a zero, bit 1.2247.7 indicates that the SC-FEC receiver has not achieved FAS lock for lane 15 of the PMA service interface This bit reflects the state of fas_lock<15> (see 153.2.4.1.1).

45.2.1.186ai.6 SC-FEC FAS lock 14 (1.2247.6)

When read as a one, bit 1.2247.6 indicates that the SC-FEC receiver has achieved FAS lock for lane 14 of the PMA service interface. When read as a zero, bit 1.2247.6 indicates that the SC-FEC receiver has not achieved FAS lock for lane 14 of the PMA service interface This bit reflects the state of fas_lock<14> (see 153.2.4.1.1).

45.2.1.186ai.7 SC-FEC FAS lock 13 (1.2247.5)

When read as a one, bit 1.2247.5 indicates that the SC-FEC receiver has achieved FAS lock for lane 13 of the PMA service interface. When read as a zero, bit 1.2247.5 indicates that the SC-FEC receiver has not achieved FAS lock for lane 13 of the PMA service interface This bit reflects the state of fas_lock<13> (see 153.2.4.1.1).

45.2.1.186ai.8 SC-FEC FAS lock 12 (1.2247.4)

When read as a one, bit 1.2247.4 indicates that the SC-FEC receiver has achieved FAS lock for lane 12 of the PMA service interface. When read as a zero, bit 1.2247.4 indicates that the SC-FEC receiver has not achieved FAS lock for lane 12 of the PMA service interface This bit reflects the state of fas_lock<12> (see 153.2.4.1.1).

45.2.1.186ai.9 SC-FEC FAS lock 11 (1.2247.3)

When read as a one, bit 1.2247.3 indicates that the SC-FEC receiver has achieved FAS lock for lane 11 of the PMA service interface. When read as a zero, bit 1.2247.3 indicates that the SC-FEC receiver has not achieved FAS lock for lane 11 of the PMA service interface This bit reflects the state of fas_lock<11> (see 153.2.4.1.1).

45.2.1.186ai.10 SC-FEC FAS lock 10 (1.2247.2)

When read as a one, bit 1.2247.2 indicates that the SC-FEC receiver has achieved FAS lock for lane 10 of the PMA service interface. When read as a zero, bit 1.2247.2 indicates that the SC-FEC receiver has not achieved FAS lock for lane 10 of the PMA service interface This bit reflects the state of fas_lock<10> (see 153.2.4.1.1).

45.2.1.186ai.11 SC-FEC FAS lock 9 (1.2247.1)

When read as a one, bit 1.2247.1 indicates that the SC-FEC receiver has achieved FAS lock for lane 9 of the PMA service interface. When read as a zero, bit 1.2247.1 indicates that the SC-FEC receiver has not achieved FAS lock for lane 9 of the PMA service interface This bit reflects the state of fas_lock<9> (see 153.2.4.1.1).

45.2.1.186ai.12 SC-FEC FAS lock 8 (1.2247.0)

When read as a one, bit 1.2247.0 indicates that the SC-FEC receiver has achieved FAS lock for lane 8 of the PMA service interface. When read as a zero, bit 1.2247.0 indicates that the SC-FEC receiver has not achieved FAS lock for lane 8 of the PMA service interface This bit reflects the state of fas_lock<8> (see 153.2.4.1.1).

45.2.1.186aj SC-FEC lane mapping, lane 0 register (Register 1.2250)

The assignment of bits in the SC-FEC lane mapping, lane 0 register is shown in Table 45–150ai. When the SC-FEC described in Clause 153 detects and locks the FAS on PMA service interface lane 0, the detected SC-FEC lane number is recorded in this register. The contents of the SC-FEC lane mapping, lane 0 register is valid when the SC-FEC FAS lock 0 bit (1.2246.0) is set to one and is invalid otherwise (see 45.2.1.186ah.9).

Table 45–150ai—SC-FEC lane mapping, lane 0 register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2250.15:5	Reserved	Value always 0	RO
1.2250.4:0	Lane 0 mapping	SC-FEC lane received on PMA service interface lane 0	RO

^aRO = Read only

45.2.1.186ak SC-FEC lane mapping, lane 1 through 19 registers (Registers 1.2251 through 1.2269)

The definition of the SC-FEC lane mapping, lane 1 through 19 registers is equivalent to that described for lane 0 in 45.2.1.186aj. The lane mapping for lane 1 is in register 1.2251; lane 2 is in register 1.2252; etc.

45.2.1.186al SC-FEC corrected codewords counter (Register 1.2276, 1.2277)

The assignment of bits in the SC-FEC corrected codewords counter is shown in Table 45–150aj. See 153.2.5.1 for a definition of this counter. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2276 and 1.2277 are used to read the value of a 32-bit counter. When registers 1.2276 and 1.2277 are used to read the 32-bit counter value, register 1.2276 is read first, the value of the register 1.2277 is latched when (and only when) register 1.2276 is read, and reads of register 1.2277 return the latched value rather than the current value of the counter.

Table 45–150aj—SC-FEC corrected codewords counter bit definitions

Bit(s)	Name	Description	R/W ^a
1.2276.15:0	FEC corrected codewords lower	FEC_corrected_cw_counter[15:0]	RO, NR
1.2277.15:0	FEC corrected codewords upper	FEC_corrected_cw_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

45.2.1.186am SC-FEC uncorrected codewords counter (Register 1.2278, 1.2279)

The assignment of bits in the SC-FEC uncorrected codewords counter is shown in Table 45–150ak. See 153.2.5.2 for a definition of this counter. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2278 and 1.2279 are used to read the value of a 32-bit counter. When registers 1.2278 and 1.2279 are used to read the 32-bit counter value, register 1.2278 is read first, the value of the register 1.2279 is latched when (and only when) register 1.2278 is read, and reads of register 1.2279 return the latched value rather than the current value of the counter.

Table 45–150ak—SC-FEC uncorrected codewords counter bit definitions

Bit(s)	Name	Description	R/W ^a
1.2278.15:0	FEC uncorrected codewords lower	FEC_uncorrected_cw_counter[15:0]	RO, NR
1.2279.15:0	FEC uncorrected codewords upper	FEC_uncorrected_cw_counter[31:16]	RO, NR

^aRO = Read only, NR = Non Roll-over

45.2.1.186an SC-FEC total bits register (Register 1.2280, 1.2281, 1.2282, 1.2283)

The assignment of bits in the SC-FEC total bits register is shown in Table 45–150al. See 153.2.5.3 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2280, 1.2281, 1.2282, and 1.2283 are used to read the value of a 64-bit counter. When registers 1.2280, 1.2281, 1.2282, and 1.2283 are used to read the 64-bit counter value, register 1.2280 is read first, the values of registers 1.2281, 1.2282, and 1.2283 are latched when (and only when) register 1.2280 is read, and reads of registers 1.2281, 1.2282, and 1.2283 return the latched value rather than the current value of the counter.

Table 45–150al—SC-FEC total bits register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2280.15:0	FEC total bits lowest	FEC_total_bits[15:0]	RO, NR
1.2281.15:0	FEC total bits	FEC_total_bits[31:16]	RO, NR
1.2282.15:0	FEC total bits	FEC_total_bits[47:32]	RO, NR
1.2283.15:0	FEC total bits highest	FEC_total_bits[64:48]	RO, NR

^aRO = Read only, NR = Non Roll-over

45.2.1.186ao SC-FEC corrected bits register (Register 1.2284, 1.2285, 1.2286, 1.2287)

The assignment of bits in the SC-FEC corrected bits register is shown in Table 45–150am. See 153.2.5.4 for a definition of this register. These bits shall be reset to all zeros when the register is read by the management function or upon PHY reset. These bits shall be held at all ones in the case of overflow. Registers 1.2284, 1.2285, 1.2286, and 1.2287 are used to read the value of a 64-bit counter. When registers 1.2284, 1.2285, 1.2286, and 1.2287 are used to read the 64-bit counter value, register 1.2284 is read first, the values of

registers 1.2285, 1.2286, and 1.2287 are latched when (and only when) register 1.2284 is read, and reads of registers 1.2285, 1.2286, and 1.2287 return the latched value rather than the current value of the counter.

Table 45–150am—SC-FEC corrected bits register bit definitions

Bit(s)	Name	Description	R/W ^a
1.2284.15:0	FEC corrected bits lowest	FEC_corrected_bits[15:0]	RO, NR
1.2285.15:0	FEC corrected bits	FEC_corrected_bits[31:16]	RO, NR
1.2286.15:0	FEC corrected bits	FEC_corrected_bits[47:32]	RO, NR
1.2287.15:0	FEC corrected bits highest	FEC_corrected_bits[64:48]	RO, NR

^aRO = Read only, NR = Non Roll-over

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

78. Energy-Efficient Ethernet (EEE)

78.1 Overview

78.1.4 PHY types optionally supporting EEE

Insert a new row for 100GBASE-ZR in Table 78–1 after 100GBASE-ER4 and before 200GBASE-KR4 (as inserted by IEEE Std 802.3cd-2018) as follows (unchanged rows not shown):

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

PHY or interface type	Clause
...	
100GBASE-ZR ^b	82, 83, 91, 135, 152, 153, 154
...	

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

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

80. Introduction to 40 Gb/s and 100 Gb/s networks

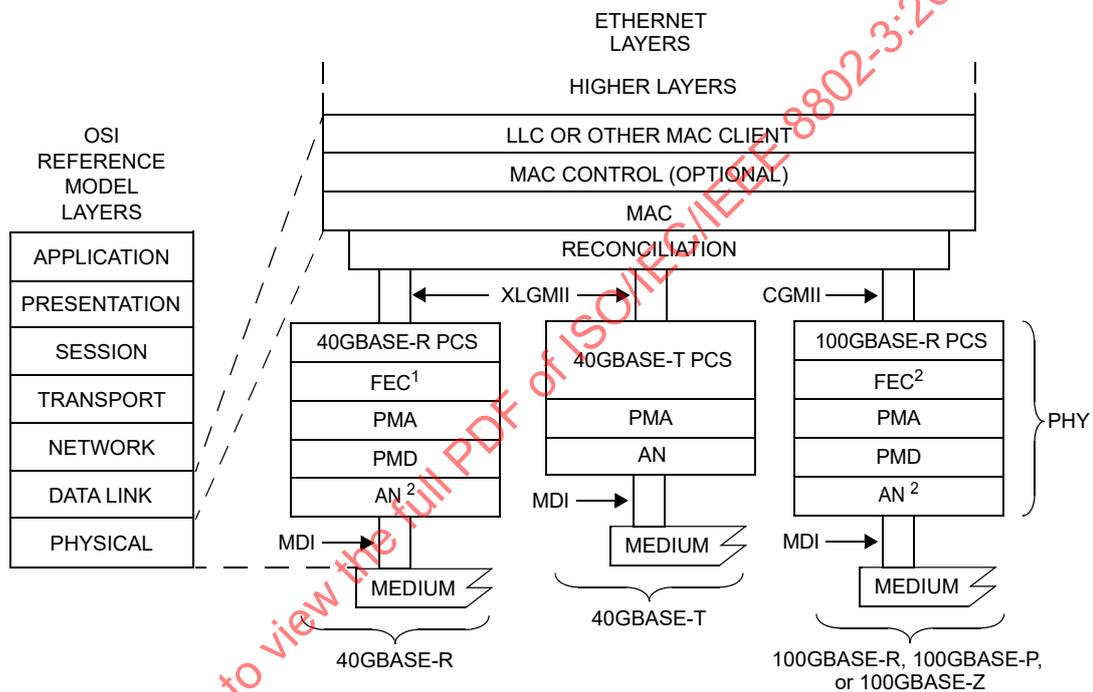
80.1 Overview

80.1.3 Relationship of 40 Gigabit and 100 Gigabit Ethernet to the ISO OSI reference model

Change list item h) in 80.1.3 (as changed by IEEE Std 802.3cu-2021) as follows:

- h) The MDIs as specified in Clause 89 for 40GBASE-FR, and in Clause 140 for 100GBASE-DR, 100GBASE-FR1, and 100GBASE-LR1, and in Clause 154 for 100GBASE-ZR use a single lane data path.

Replace Figure 80–1 with the following figure:



AN = AUTO-NEGOTIATION
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 XLGMII = 40 Gb/s MEDIA INDEPENDENT INTERFACE
 NOTE 1—OPTIONAL OR OMITTED DEPENDING ON PHY TYPE
 NOTE 2—CONDITIONAL BASED ON PHY TYPE

Figure 80–1—Architectural positioning of 40 Gigabit and 100 Gigabit Ethernet

80.1.4 Nomenclature

Insert a new paragraph into 80.1.4 before the seventh paragraph beginning “Physical Layer devices ...” as follows:

100GBASE-Z represents a family of Physical Layer devices using the Clause 82 Physical Coding Sublayer for 100 Gb/s operation over multiple PCS lanes (see Clause 82), the FEC of Clause 153, and a PMD

implementing dual polarization differential quadrature phase shift keying (DP-DQPSK) modulation. Some 100GBASE-Z Physical Layer devices also use the FEC of Clause 91 and the Inverse RS-FEC of Clause 152.

Insert a new row at the end of Table 80–1 as follows (unchanged rows not shown):

Table 80–1—40 Gb/s and 100 Gb/s PHYs

Name	Description
...	
100GBASE-ZR	100 Gb/s PHY using 100GBASE-R encoding capable of transmission over a specified channel on a defined DWDM grid in each direction of transmission with reach up to at least 80 km (see Clause 154)

80.1.5 Physical Layer signaling systems

Insert Table 80–4b after Table 80–4a (as inserted by IEEE Std 802.3cd-2018) as follows:

Table 80–4b—Nomenclature and clause correlation (100GBASE-Z optical)

Nomenclature	Clause ^a																	
	78	81	82	83	83A	83B	83D	83E	91	135	135D	135E	135F	135G	152	153	154	
	EEE	RS	CGMII	100GBASE-R PCS	100GBASE-R PMA	CAUI-10	CAUI-10	CAUI-4	CAUI-4	RS-FEC	100GBASE-P PMA	100GAUI-4 C2C	100GAUI-4 C2M	100GAUI-2 C2C	100GAUI-2 C2M	Inverse RS-FEC	SC-FEC and 100GBASE-ZR PMA	100GBASE-ZR PMD
100GBASE-ZR	O	M	O	M	O	O	O	O	O	O	O	O	O	O	O	C ^b	M	M

^a O = Optional, M = Mandatory, C = Conditional.

^b Clause 152 inverse RS-FEC mandatory when Clause 91 RS-FEC is present.

80.2 Summary of 40 Gigabit and 100 Gigabit Ethernet sublayers

80.2.2 Physical Coding Sublayer (PCS)

Change 80.2.2 (as changed by IEEE Std 802.3cd-2018) as follows:

The terms 40GBASE-R, 100GBASE-R, ~~and 100GBASE-P~~, and 100GBASE-Z refer to a specific family of Physical Layer implementations based upon the 64B/66B data coding method specified in Clause 82 and the PMA specifications defined in Clause 83, Clause 94, ~~or Clause 135~~, or Clause 153. Clause 82 PCSs perform

encoding (decoding) of data from (to) the XLGMII/CGMII to 64B/66B code blocks, distribute the data to multiple lanes, and transfer the encoded data to the PMA or FEC.

80.2.3 Forward Error Correction (FEC) sublayers

Change 80.2.3 (as changed by IEEE Std 802.3cu-2021) as follows:

A Forward Error Correction sublayer is optional for 40GBASE-KR4, 40GBASE-CR4, and 100GBASE-CR10 PHYs and mandatory for 100GBASE-CR2, 100GBASE-CR4, 100GBASE-KR2, 100GBASE-KR4, 100GBASE-KP4, 100GBASE-SR2, 100GBASE-SR4, 100GBASE-DR, 100GBASE-FR1, ~~and~~ 100GBASE-LR1, and 100GBASE-ZR PHYs. The FEC sublayer can be placed in between the PCS and PMA sublayers or between two PMA sublayers.

The BASE-R FEC (see [Clause 74](#)) is instantiated for each PCS lane and operates autonomously on a per PCS lane basis. The Reed-Solomon FEC (see [Clause 91](#)) is instantiated once and requires 20 PCS lanes and 4 PMA lanes for operation. The SC-FEC (see [Clause 153](#)) is instantiated once and requires 20 PCS lanes and 20 PMA lanes for operation.

80.2.4 Physical Medium Attachment (PMA) sublayer

Change the second paragraph of 80.2.4 (as changed by IEEE Std 802.3cd-2018) as follows:

~~The 40GBASE-R and 100GBASE-R PMAs are specified in [Clause 83](#). The PMA specific to the 100GBASE-KP4 PHY is specified in [Clause 94](#). The PMA for other 100GBASE-P PHYs is specified in [Clause 135](#).~~ [Clause 83](#) specifies 40GBASE-R and 100GBASE-R PMAs that may be used with any PHY type of the corresponding rate. Additional PMAs are only applicable to specific PHY types:

- a) [Clause 94](#) specifies a PMA that may be used only in a 100GBASE-KP4 PHY.
- b) [Clause 135](#) specifies a PMA that may be used in other 100GBASE-P or 100GBASE-ZR PHY types.
- c) [Clause 153](#) specifies a PMA that is used in the 100GBASE-ZR PHY.

80.2.5 Physical Medium Dependent (PMD) sublayer

Change the second paragraph of 80.2.5 (as changed by IEEE Std 802.3cd-2018) as follows:

The 40GBASE-R, 100GBASE-R, and 100GBASE-P PMDs and their corresponding media are specified in [Clause 84](#) through [Clause 89](#), [Clause 92](#) through [Clause 95](#), [Clause 136](#) through [Clause 138](#), ~~and [Clause 140](#), and [Clause 154](#).~~

80.3 Service interface specification method and notation

80.3.2 Instances of the Inter-sublayer service interface

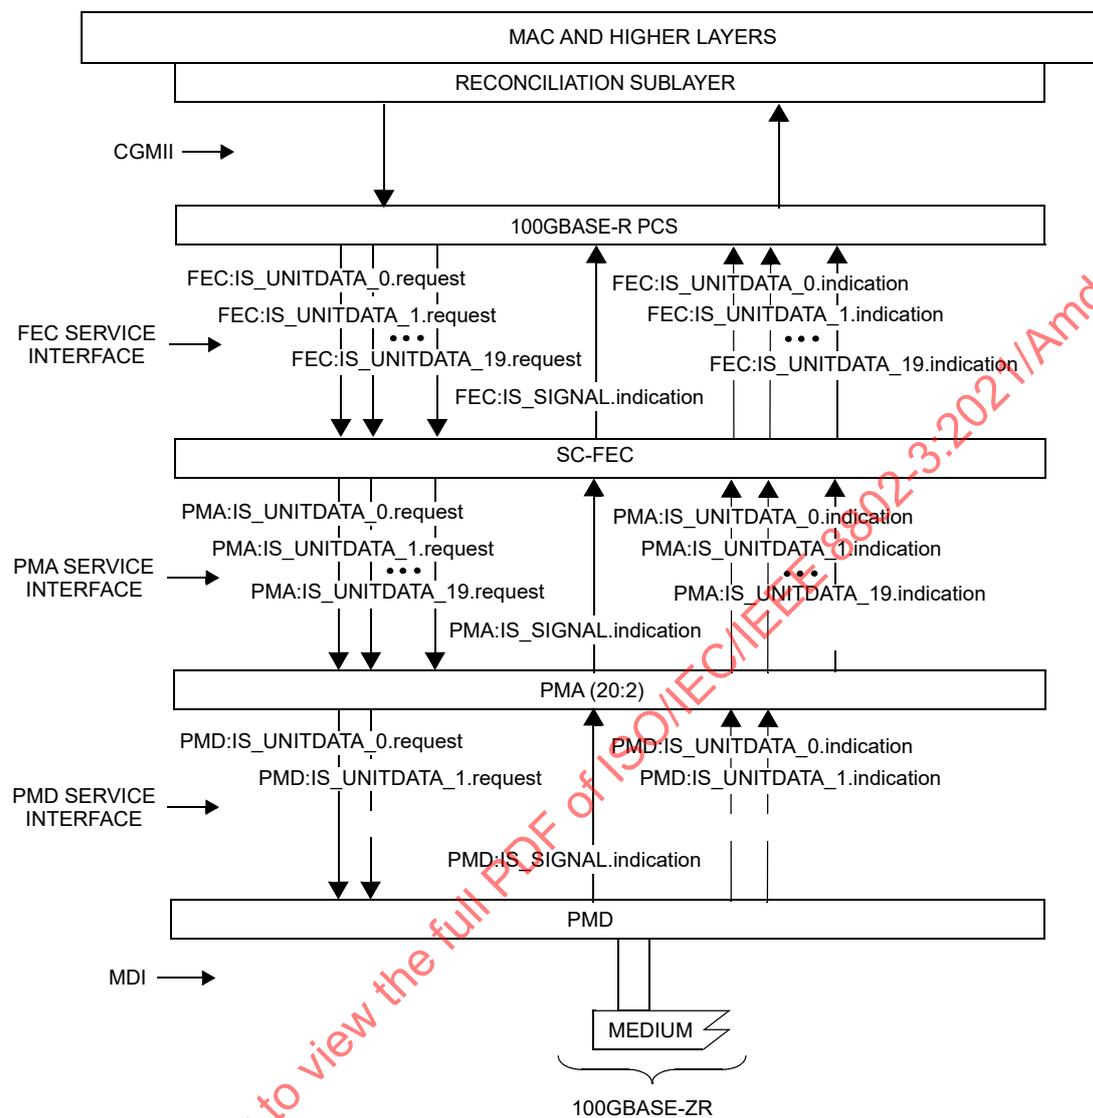
Insert item d) at the end of the list in 80.3.2 as follows:

- d) IFEC:—for primitives issued on the interface between the Inverse FEC sublayer and the PMA sublayer called the Inverse FEC service interface.

Change the first sentence of the second paragraph of 80.3.2 as follows:

Examples of inter-sublayer service interfaces for 40GBASE-R, 100GBASE-R, ~~and~~ 100GBASE-P, and 100GBASE-Z with their corresponding instance names are illustrated in [Figure 80–2](#), [Figure 80–3](#), [Figure 80–4](#), [Figure 80–4a](#), and [Figure 80–5](#).

Insert new Figure 80-4a after Figure 80-4 as follows:



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE

PCS = PHYSICAL CODING SUBLAYER
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 SC-FEC = STAIRCASE FORWARD ERROR CORRECTION

Figure 80-4a—100GBASE-Z inter-sublayer service interfaces with SC-FEC

80.4 Delay constraints

Insert new rows into Table 80–5 with the rows for sublayers “Inverse RS-FEC” and “SC-FEC” after the row for “100GBASE-R RS-FEC” and the row for “100GBASE-ZR PMD” at the bottom of the table as follows (unchanged rows not shown):

Table 80–5—Sublayer delay constraints

Sublayer	Maximum (bit time) ^a	Maximum (pause_quanta) ^b	Maximum (ns)	Notes ^c
...				
Inverse RS-FEC	40 960	80	409.60	See 152.4.
SC-FEC	1 827 840	3570	18 278.40	See 153.2.2.
...				
100GBASE-ZR PMD	2048	4	20.48	Includes 2 m of fiber. See 154.3.1.

^a For 40GBASE-R, 1 bit time (BT) is equal to 25 ps and for 100GBASE-R, 1 bit time (BT) is equal to 10 ps. (See 1.4.160 for the definition of bit time.)
^b For 40GBASE-R, 1 pause_quantum is equal to 12.8 ns and for 100GBASE-R, 1 pause_quantum is equal to 5.12 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.

80.5 Skew constraints

Change Table 80–6 (as modified by IEEE Std 802.3cd-2018) as follows:

Table 80–6—Summary of Skew constraints

Skew points	Maximum Skew (ns) ^a	Maximum Skew for 40GBASE-R PCS lane (UI) ^b	Maximum Skew for 100GBASE-R PCS lane (UI) ^c	Notes ^d
SP0	29	N/A	≈ 150	See 83.5.3.1 or 135.5.3
SP1	29	≈ 299	≈ 150	See 83.5.3.2 or 135.5.3
SP2	43	≈ 443	≈ 222	See 83.5.3.4, 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, 95.3.2, 135.5.3, 135.5.3, 136.6, 137.6, 138.3, or 140.3 , or 154.3.2
SP3	54	≈ 557	≈ 278	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, 95.3.2, 135.5.3, 136.6, 137.6, 138.3, or 140.3 , or 154.3.2
SP4	134	≈ 1382	≈ 691	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, 95.3.2, 135.5.3, 136.6, 137.6, 138.3, or 140.3 , or 154.3.2

Table 80–6—Summary of Skew constraints (continued)

Skew points	Maximum Skew (ns) ^a	Maximum Skew for 40GBASE-R PCS lane (UI) ^b	Maximum Skew for 100GBASE-R PCS lane (UI) ^c	Notes ^d
SP5	145	≈ 1495	≈ 748	See 84.5, 85.5, 86.3.2, 87.3.2, 88.3.2, 89.3.2, 92.5, 93.5, 94.3.4, 95.3.2, 135.5.3, 136.6, 137.6, 138.3, or 140.3, or 154.3.2
SP6	160	≈ 1649	≈ 824	See 83.5.3.6 or 135.5.3
SP7	29	N/A	≈ 150	See 83.5.3.8 or 135.5.3
At PCS receive	180	≈ 1856	≈ 928	See 82.2.13
At RS-FEC transmit	49	N/A	≈ 253	See 91.5.2.2
At RS-FEC receive ^e	180	N/A	≈ 4641	See 91.5.3.1
At PCS receive (with RS-FEC)	49	N/A	≈ 253	See 82.2.13

^aThe Skew limit includes 1 ns allowance for PCB traces that are associated with the Skew points.

^bThe symbol ≈ indicates approximate equivalent of maximum Skew in UI for 40GBASE-R, based on 1 UI equals 96.969697 ps at PCS lane signaling rate of 10.3125 GBd.

^cThe symbol ≈ indicates approximate equivalent of maximum Skew in UI for 100GBASE-R, based on 1 UI equals 193.939394 ps at PCS lane signaling rate of 5.15625 GBd.

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

^eThe skew at the RS-FEC receive is the skew between RS-FEC lanes. The symbol ≈ indicates approximate equivalent of maximum Skew in UI for RS-FEC lanes with a signaling rate of 25.78125 GBd.

80.7 Protocol implementation conformance statement (PICS) proforma

Change the first paragraph of 80.7 (as modified by IEEE Std 802.3cd-2018) as follows:

The supplier of a protocol implementation that is claimed to conform to any part of IEEE Std 802.3, Clause 45, Clause 73, Clause 74, Clause 81 through Clause 89, Clause 91 through Clause 95, Clause 135 through Clause 138, Clause 140, Clause 152 through Clause 154, and related annexes demonstrates compliance by completing a protocol implementation conformance statement (PICS) proforma.

82. Physical Coding Sublayer (PCS) for 64B/66B, type 40GBASE-R and 100GBASE-R

82.2 Physical Coding Sublayer (PCS)

82.2.3 64B/66B transmission code

82.2.3.3 Block structure

Change the warning in 80.2.3.3 as follows:

WARNING

The mapping of 40GBASE-R PCS blocks into OPU3 specified in ITU-T G.709 [B48] depends on the set of control block types shown in Figure 82–5. Any deviation from the coding specified in Figure 82–5 will break the mapping and may prevent 40GBASE-R PCS blocks from being mapped into OPU3 (see ITU-T G.709 [B48] for more details).

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

Insert new Clause 152, Clause 153, and Clause 154 after Clause 151 as follows:

152. Inverse RS-FEC sublayer

152.1 Overview

152.1.1 Scope

The Inverse RS-FEC sublayer specifies a Reed-Solomon forward error correction (RS-FEC) sublayer for 100GBASE-R, 100GBASE-P, and 100GBASE-Z PHYs. This sublayer is used in cases where the RS-FEC specified in Clause 91 is used across a physically instantiated 100GAUI-n and a different FEC is used for the PMD.

152.1.2 Position of Inverse RS-FEC in the 100GBASE-R sublayers

Figure 152–1 shows the relationship of the Inverse RS-FEC sublayer to the ISO/IEC Open System Interconnection (OSI) reference model.

152.2 Inverse RS-FEC service interface

This subclause specifies the services provided by the Inverse RS-FEC sublayer. The service interface is described in an abstract manner and does not imply any particular implementation. The Inverse RS-FEC service interface is an instance of the inter-sublayer service interface defined in 80.3. The service interface primitives are summarized as follows:

```
IFEC:IS_UNITDATA_ i.request
IFEC:IS_UNITDATA_ i.indication
IFEC:IS_SIGNAL.indication
```

The primitives are defined for $i=0$ to 3. The PMA continuously sends 4 parallel bit streams to the Inverse RS-FEC sublayer, each at a nominal signaling rate of 26.5625 GBd. The Inverse RS-FEC sublayer continuously sends 4 parallel bit streams to the PMA, each at a nominal signaling rate of 26.5625 GBd.

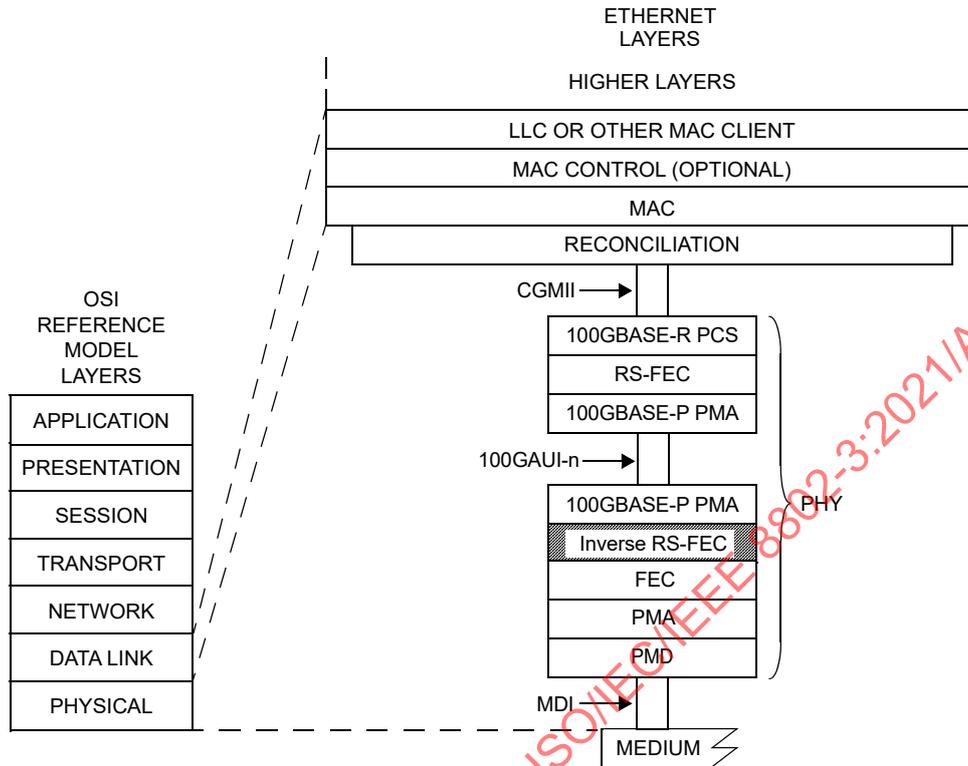
The SIGNAL_OK parameter of the IFEC:IS_SIGNAL.indication primitive can take one of two values: OK or FAIL. The value is set to OK when align_status (see 152.6.12) is true. The value is set to FAIL when align_status is false.

152.3 PMA or FEC sublayer compatibility

The Inverse RS-FEC sublayer requires that the PMA or FEC service interface below the Inverse RS-FEC sublayer consists of twenty upstream lanes and twenty downstream lanes. The restriction that all PMA service interfaces between the RS-FEC sublayer and the PMD sublayer consist of four or fewer lanes (see 91.3) is removed below the Inverse RS-FEC sublayer.

152.4 Delay constraints

The maximum delay contributed by the Inverse RS-FEC sublayer (sum of transmit and receive delays at one end of the link) shall be no more than 40 960 bit times (80 pause_quanta or 409.6 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.



100GAUI-n = 100 Gb/s n-LANE ATTACHMENT UNIT INTERFACE
 CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 FEC = FORWARD ERROR CORRECTION
 Inverse RS-FEC = Inverse REED-SOLOMON FORWARD ERROR CORRECTION
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 RS-FEC = REED-SOLOMON FORWARD ERROR CORRECTION

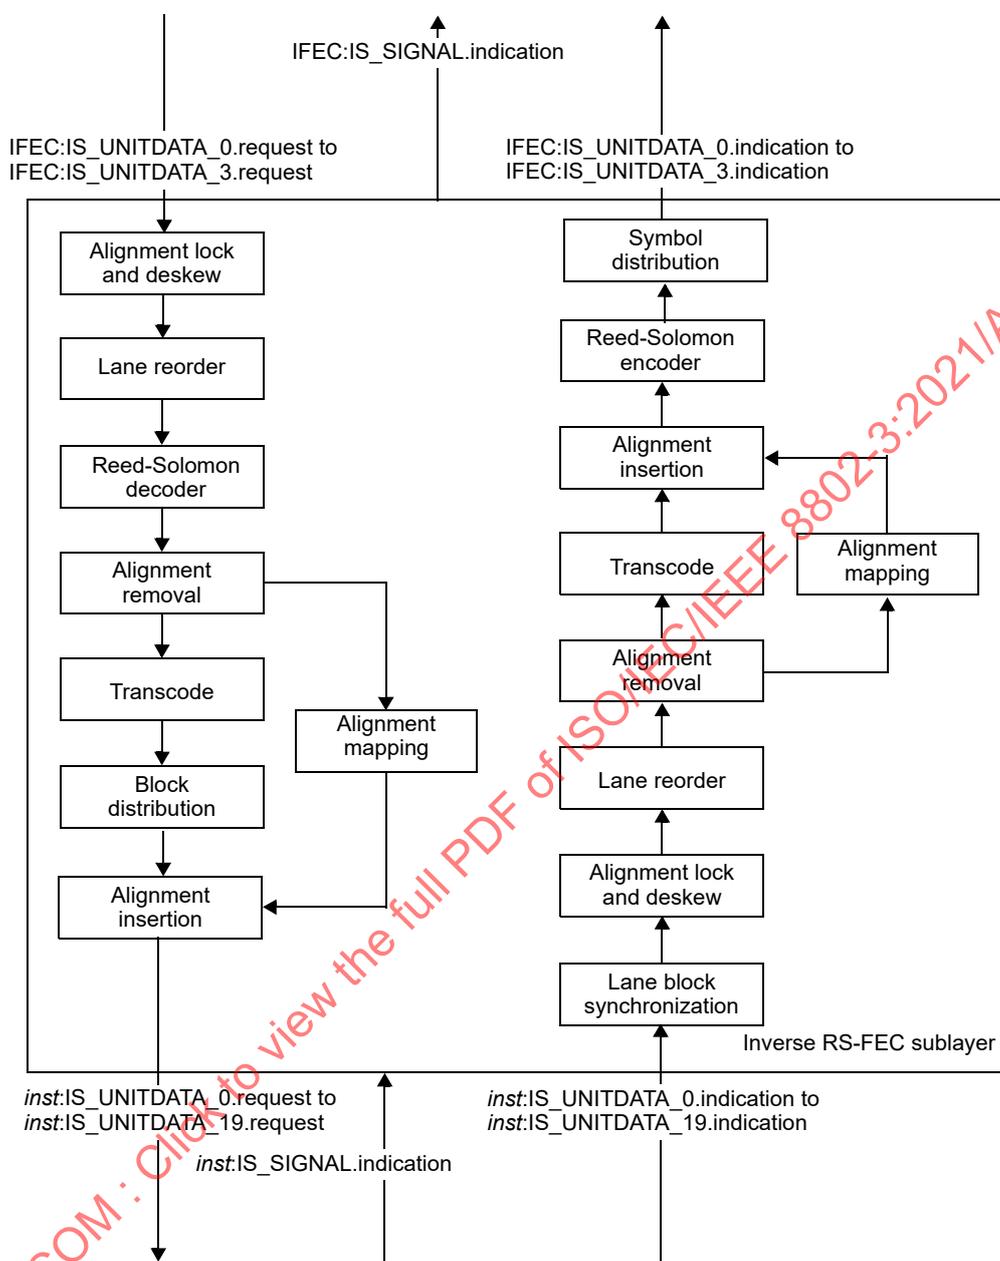
Figure 152-1—Inverse RS-FEC relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

152.5 Functions within the Inverse RS-FEC sublayer

The Inverse RS-FEC sublayer performs a subset of the functions defined for the RS-FEC sublayer described in Clause 91 with the Tx and Rx directions reversed. Specific limitations as compared to Clause 91 are that only RS(544,514) FEC is supported and EEE deep sleep mode is not supported. The FEC optional states in Clause 91 are mandatory for the Inverse RS-FEC sublayer.

152.5.1 Functional block diagram

A functional block diagram of the Inverse RS-FEC sublayer is shown in Figure 152-2.



inst PMA or FEC, depending on which sublayer is below the Inverse RS-FEC sublayer

Figure 152–2—Inverse RS-FEC sublayer functional block diagram

152.5.2 Transmit function

152.5.2.1 Alignment lock and deskew

The Inverse RS-FEC transmit function forms 4 bit streams by concatenating the bits received from each of the 4 IFEC:IS_UNITDATA_0.request primitives in the order they are received. It obtains lock to the

alignment markers as specified by the FEC synchronization state diagram shown in Figure 91–8. The FEC optional states and transition A in Figure 91–8 are mandatory for the Inverse RS-FEC sublayer.

After alignment marker lock is achieved on all 4 lanes, all inter-lane Skew is removed as specified by the FEC alignment state diagram shown in Figure 91–9. The Inverse RS-FEC transmit function shall support a maximum Skew of 49 ns between FEC lanes and a maximum Skew Variation of 400 ps.

152.5.2.2 Lane reorder

A particular FEC lane can be received on any of the lanes of the Inverse FEC service interface. The Inverse RS-FEC transmit function shall order the FEC lanes according to the FEC lane number (see 91.5.2.6). The FEC lane number is defined by the sequence of alignment markers that are mapped to each FEC lane.

After all FEC lanes are aligned, deskewed, and reordered, the FEC lanes are multiplexed together in the proper order to reconstruct the original stream of FEC codewords.

152.5.2.3 Reed-Solomon decoder

The Reed-Solomon decoder implements the RS(544,514) FEC decoder described in 91.5.3.3 with the exception that message symbols come from tx_scrambled rather than rx_scrambled. The optional subclause 91.5.3.3.1 is not supported for the Inverse RS-FEC sublayer.

152.5.2.4 Alignment marker removal

The first 1285 message bits in every 4096th codeword is the vector am_txmapped<1284:0> where bit 0 is the first bit received. The specific codewords that include this vector are indicated by the alignment lock and deskew function (refer to 152.5.2.1).

The vector am_txmapped shall be removed prior to transcoding.

152.5.2.5 256B/257B to 64B/66B transcoder

The transcoder extracts a group of four 66-bit blocks, tx_coded_j<65:0> where j = 0 to 3, from each 257-bit block tx_scrambled<256:0>. Bit 0 of the 257-bit block is the first bit received.

First, descramble the first 5 bits, based on reception order, of tx_scrambled<256:0> to yield tx_xcoded<256:0> as follows:

- a) Set tx_xcoded<4:0> to the result of the bit-wise exclusive-OR of tx_scrambled<4:0> and tx_scrambled<12:8>.
- b) Set tx_xcoded<256:5> to tx_scrambled<256:5>.

If tx_xcoded<0> is 1, tx_coded_j<65:0> for j = 0 to 3 shall be derived as follows:

- a1) tx_coded_j<65:2> = tx_xcoded<(64j + 64):(64j + 1)> for j = 0 to 3.
- b1) tx_coded_j<0> = 0 and tx_coded_j<1> = 1 for all j = 0 to 3.

If tx_xcoded<0> is 0 and any tx_xcoded<j + 1> = 0 for j = 0 to 3, tx_coded_j<65:0> for j = 0 to 3 shall be derived as follows:

- a2) Let c be the smallest value of j such that tx_xcoded<j + 1> = 0. In other words, tx_coded_c is the first 66-bit control block in the resulting group of four blocks.

- b2) Let tx_payloads be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions:
- $$\text{tx_payloads}\langle(64c + 3):0\rangle = \text{tx_xcoded}\langle(64c + 8):5\rangle$$
- $$\text{tx_payloads}\langle(64c + 7):(64c + 4)\rangle = 0000 \text{ (an arbitrary value that is later replaced by } s_c)$$
- $$\text{tx_payloads}\langle 255:(64c + 8)\rangle = \text{tx_xcoded}\langle 256:(64c + 9)\rangle$$
- c2) $\text{tx_coded}_j\langle 65:2\rangle = \text{tx_payloads}\langle(64j + 63):64j\rangle$ for $j = 0$ to 3.
- d2) Let $f_c\langle 3:0\rangle = \text{tx_coded}_c\langle 5:2\rangle$ be the scrambled first nibble (based on transmission order) of the block type field for tx_coded_c .
- e2) Descramble $f_c\langle 3:0\rangle$ to yield $g\langle 3:0\rangle$ per the following expression where “^” denotes the exclusive-OR operation. When $c = 0$, $\text{tx_coded}_{(c-1)}$ corresponds to tx_coded_3 from the previous 257-bit block.
- $$g\langle i\rangle = f_c\langle i\rangle \wedge \text{tx_coded}_{(c-1)}\langle i + 8\rangle \wedge \text{tx_coded}_{(c-1)}\langle i + 27\rangle \text{ for } i = 0 \text{ to } 3.$$
- f2) The block type field may be uniquely identified by either its most or least significant nibble. Since $g\langle 3:0\rangle$ is the least significant nibble of the block type field (per the transmission order), derive $h\langle 3:0\rangle$ by cross-referencing to $g\langle 3:0\rangle$ using Figure 82–5. For example, if $g\langle 3:0\rangle$ is 0xE then $h\langle 3:0\rangle$ is 0x1. If no match to $g\langle 3:0\rangle$ is found, $h\langle 3:0\rangle$ is set to 0000.
- g2) If $\text{tx_xcoded}\langle j + 1\rangle = 0$, $\text{tx_coded}_j\langle 0\rangle = 1$ and $\text{tx_coded}_j\langle 1\rangle = 0$ for $j = 0$ to 3.
- h2) If $\text{tx_xcoded}\langle j + 1\rangle = 1$, $\text{tx_coded}_j\langle 0\rangle = 0$ and $\text{tx_coded}_j\langle 1\rangle = 1$ for $j = 0$ to 3.
- i2) If $h\langle 3:0\rangle = 0000$, $\text{tx_coded}_c\langle 1\rangle = 1$ (invalidate synchronization header).

If $\text{tx_xcoded}\langle 0\rangle$ is 0 and all $\text{tx_xcoded}\langle j + 1\rangle = 1$ for $j = 0$ to 3, $\text{tx_coded}_j\langle 65:0\rangle$ for $j = 0$ to 3 shall be derived as follows:

- a3) Set $c = 0$ and $h\langle 3:0\rangle = 0000$.
- b3) Let tx_payloads be a vector representing the payloads of the four 66-bit blocks. It is derived using the following expressions.
- $$\text{tx_payloads}\langle(64c + 3):0\rangle = \text{tx_xcoded}\langle(64c + 8):5\rangle$$
- $$\text{tx_payloads}\langle(64c + 7):(64c + 4)\rangle = 0000 \text{ (an arbitrary value that is later replaced by } s_c)$$
- $$\text{tx_payloads}\langle 255:(64c + 8)\rangle = \text{tx_xcoded}\langle 256:(64c + 9)\rangle$$
- c3) $\text{tx_coded}_j\langle 65:2\rangle = \text{tx_payloads}\langle(64j + 63):(64j)\rangle$ for $j = 0$ to 3.
- d3) $\text{tx_coded}_j\langle 0\rangle = 0$ and $\text{tx_coded}_j\langle 1\rangle = 0$ for $j = 0$ and 2.
- e3) $\text{tx_coded}_j\langle 0\rangle = 1$ and $\text{tx_coded}_j\langle 1\rangle = 1$ for $j = 1$ and 3.

If $\text{tx_xcoded}\langle 0\rangle$ is 0, scramble $h\langle 3:0\rangle$ to yield $s_c\langle 3:0\rangle$ and assign it to tx_coded_c per the following expressions.

- a4) $s_c\langle i\rangle = h\langle i\rangle \wedge \text{tx_coded}_{(c-1)}\langle i + 12\rangle \wedge \text{tx_coded}_{(c-1)}\langle i + 31\rangle$ for $i = 0$ to 3.
- b4) $\text{tx_coded}_c\langle 9:6\rangle = s_c\langle 3:0\rangle$.

The 66-bit blocks are transmitted in order from $j = 0$ to 3. Bit 0 of each block is the first bit transmitted.

152.5.2.6 Block distribution

After the data has been transcoded, it shall be distributed to 20 PCS lanes, one 66-bit block at a time in a round-robin distribution from the lowest to the highest numbered PCS lanes. The distribution process is shown in Figure 82–6.

152.5.2.7 Alignment marker mapping and insertion

The alignment marker mapping function compensates for the operation of the lane reorder function (refer to 152.5.2.2) to derive the PCS lane alignment markers, $\text{am_tx}_x\langle 65:0\rangle$ for $x = 0$ to 19, from $\text{am_txmapped}\langle 1284:0\rangle$ (refer to 152.5.2.4).

The alignment markers shall be derived from $\text{am_txmapped}\langle 1284:0 \rangle$ in a manner that yields the same result as the following process.

Given $i = 0$ to 3, $k = 0$ to 31, and $y = i + 4k$, am_txpayloads may be derived from am_txmapped per the following expression:

$$\text{am_txpayloads}\langle i, (10k + 9):10k \rangle = \text{am_txmapped}\langle (10y + 9):10y \rangle$$

The 5-bit pad $\text{am_txmapped}\langle 1284:1280 \rangle$ is ignored. Given $i = 0$ to 3, $j = 0$ to 4, and $x = i + 4j$, amp_tx_x may be derived from am_txpayloads by the following expression:

$$\text{amp_tx_x}\langle 63:0 \rangle = \text{am_txpayloads}\langle i, (64j + 63):64j \rangle$$

For $x = 0$ to 19, $\text{am_tx_x}\langle 65:0 \rangle$ is constructed as follows:

- a) $\text{am_tx_x}\langle 0 \rangle = 1$ and $\text{am_tx_x}\langle 1 \rangle = 0$.
- b) $\text{am_tx_x}\langle 25:2 \rangle$ is set to M_0 , M_1 , and M_2 as shown in Figure 82–9 using the values in Table 82–2 for PCS lane number x .
- c) $\text{am_tx_x}\langle 33:26 \rangle = \text{amp_tx_x}\langle 31:24 \rangle$.
- d) $\text{am_tx_x}\langle 57:34 \rangle$ is set to M_4 , M_5 , and M_6 as shown in Figure 82–9 using the values in Table 82–2 for PCS lane number x .
- e) $\text{am_tx_x}\langle 65:58 \rangle = \text{amp_tx_x}\langle 63:56 \rangle$.

One vector is mapped to 20 alignment markers every 4096 Reed-Solomon codewords (see 152.5.2.4). The alignment markers are simultaneously transmitted on the 20 PCS lanes after every 16 383rd column of 20 66-bit blocks.

The alignment markers am_tx_0 to am_tx_3 shall be inserted so that they are immediately followed by tx_coded_0 to tx_coded_3 , respectively, as derived from the first 257-bit block following am_txmapped . Similarly am_tx_4 to am_tx_7 are followed by the 66-bit blocks corresponding to the second 257-bit block following am_txmapped , and so on.

152.5.2.8 Transmit bit ordering

The transmit bit ordering is illustrated in Figure 152–3. This illustration shows the case where the FEC lanes appear across the $\text{IFEC:IS_UNITDATA}_i$ request primitives in the correct order.

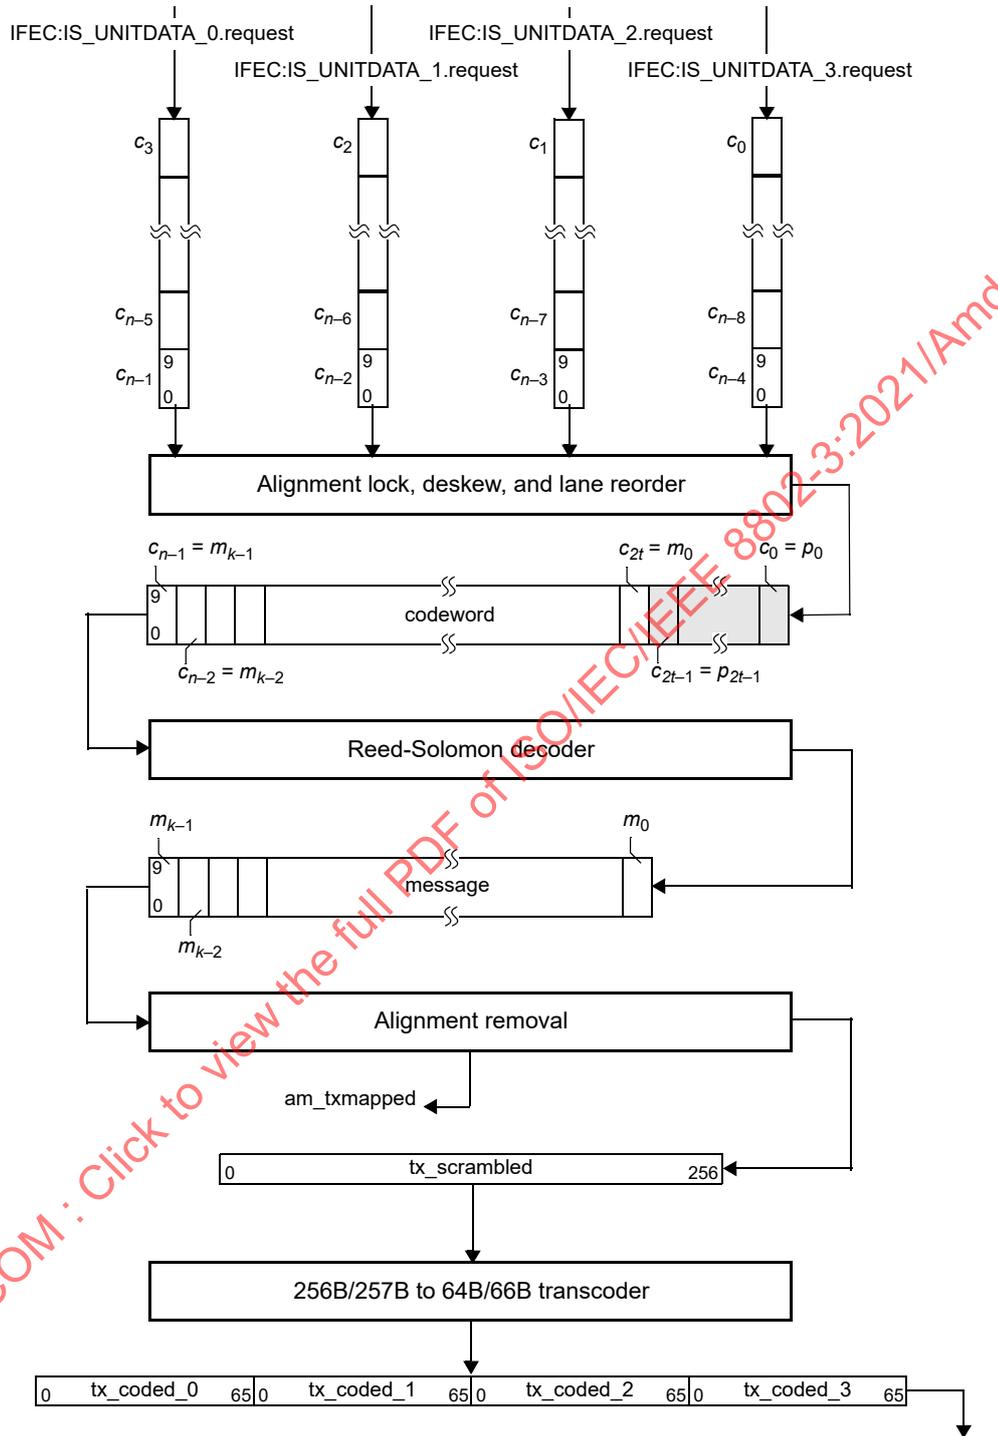


Figure 152-3—Transmit bit ordering

To Block distribution

152.5.3 Receive function

152.5.3.1 Lane block synchronization

The Inverse RS-FEC receive function forms 20 bit streams by concatenating the bits from each of the 20 FEC:IS_UNITDATA_*i*.indication primitives from the sublayer below in the order they are received. It obtains lock to the 66-bit blocks in each bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the block lock state diagram shown in Figure 82–12.

152.5.3.2 Alignment lock and deskew

Once the Inverse RS-FEC receive function achieves block lock on a PCS lane, it then begins obtaining alignment marker lock as specified in the alignment marker lock state diagram shown in Figure 82–13. This process identifies the PCS lane number received on a particular lane of the service interface of the sublayer below. After alignment marker lock is achieved on all 20 lanes, all inter-lane Skew is removed as specified by the PCS deskew state diagram shown in Figure 82–14. The Inverse RS-FEC receive function shall support a maximum Skew of 49 ns between PCS lanes and a maximum Skew Variation of 400 ps. Skew and Skew Variation are defined in 80.5.

152.5.3.3 Lane reorder

Any PCS lane may be received on any lane of the service interface. The Inverse RS-FEC receive function shall order the PCS lanes according to the PCS lane number.

152.5.3.4 Alignment marker removal

After all PCS lanes are aligned and deskewed, the PCS lanes are multiplexed together in the proper order to reconstruct the original stream of blocks and the alignment markers are removed from the data stream. Note that an alignment marker is always removed when am_lock is true for a given PCS lane even if it does not match the expected alignment marker value (due to a bit error, for example). Repeated alignment marker errors result in am_lock being set to false for a given PCS lane, but until that happens, it is sufficient to remove the block in the alignment marker position.

As part of the alignment marker removal process, the BIP₃ field is compared to the calculated Bit Interleaved Parity (BIP) value (see 82.2.8) for each PCS lane. If a Clause 45 MDIO is implemented, then the appropriate BIP error counter register (registers 1.230 to 1.249) is incremented by one each time the calculated BIP value does not equal the value received in the BIP₃ field. The bit error ratio in the data received from the far-end PCS can be estimated by dividing the BIP block error ratio by a factor of 1 081 344.

152.5.3.5 64B/66B to 256B/257B transcoder

The transcoder constructs a 257-bit block, rx_scrambled<256:0>, from a group of four 66-bit blocks, rx_coded_*j*<65:0> where *j* = 0 to 3. For each group of four 66-bit blocks, *j* = 3 corresponds to the most recently received block. Bit 0 in each 66-bit block is the first bit received and corresponds to the first bit of the synchronization header.

If for all *j* = 0 to 3, rx_coded_*j*<0> = 0 and rx_coded_*j*<1> = 1, rx_xcoded<256:0> shall be constructed as follows:

- a) rx_xcoded<0> = 1.
- b) rx_xcoded<(64*j* + 64):(64*j* + 1)> = rx_coded_<*j*:65:2> for *j* = 0 to 3.

If for all $j = 0$ to 3, $\text{rx_coded_j}<0> \neq \text{rx_coded_j}<1>$ (valid synchronization header) and for any $j = 0$ to 3, $\text{rx_coded_j}<0> = 1$ and $\text{rx_coded_j}<1> = 0$, $\text{rx_xcoded}<256:0>$ shall be constructed as follows:

- a1) $\text{rx_xcoded}<0> = 0$.
- b1) $\text{rx_xcoded}<j + 1> = \text{rx_coded_j}<1>$ for $j = 0$ to 3.
- c1) Let c be the smallest value of j such that $\text{rx_coded_c}<0> = 1$. In other words, rx_coded_c is the first 66-bit control block that was received in the current group of four blocks.
- d1) Let $\text{rx_payloads}<(64j + 63):64j> = \text{rx_coded_j}<65:2>$ for $j = 0$ to 3.
- e1) Omit $\text{rx_coded_c}<9:6>$, which is the second nibble (based on transmission order) of the block type field for rx_coded_c , from rx_xcoded per the following expressions:

$$\text{rx_xcoded}<(64c + 8):5> = \text{rx_payloads}<(64c + 3):0>$$

$$\text{rx_xcoded}<256:(64c + 9)> = \text{rx_payloads}<255:(64c + 8)>$$

If for any $j = 0$ to 3, $\text{rx_coded_j}<0> = \text{rx_coded_j}<1>$ (invalid synchronization header), $\text{rx_xcoded}<256:0>$ shall be constructed as follows:

- a2) $\text{rx_xcoded}<0> = 0$.
- b2) $\text{rx_xcoded}<j + 1> = 1$ for $j = 0$ to 3.
- c2) Let $\text{rx_payloads}<(64j + 63):64j> = \text{rx_coded_j}<65:2>$ for $j = 0$ to 3.
- d2) Omit the second nibble (based on transmission order) of rx_coded_0 per the following expressions.

$$\text{rx_xcoded}<8:5> = \text{rx_payloads}<3:0>$$

$$\text{rx_xcoded}<256:9> = \text{rx_payloads}<255:8>$$

Finally, scramble the first 5 bits, based on transmission order, of $\text{rx_xcoded}<256:0>$ to yield $\text{rx_scrambled}<256:0>$ as follows:

- a3) Set $\text{rx_scrambled}<4:0>$ to the result of the bit-wise exclusive-OR of $\text{rx_xcoded}<4:0>$ and $\text{rx_xcoded}<12:8>$.
- b3) Set $\text{rx_scrambled}<256:5>$ to $\text{rx_xcoded}<256:5>$.

For each 257-bit block, bit 0 shall be the first bit transmitted.

152.5.3.6 Alignment marker mapping and insertion

The alignment markers that were removed per 152.5.3.4 are re-inserted after being processed by the alignment marker mapping function. The alignment marker mapping function compensates for the operation of the symbol distribution function defined in 152.5.3.8 and rearranges the alignment marker bits so that they appear on the FEC lanes intact and in the desired sequence. This preserves the properties of the alignment markers (e.g., DC balance, transition density) and provides a deterministic pattern for the purpose of synchronization. The receive function in the RS-FEC sublayer above uses knowledge of this mapping to determine the FEC lane that is received on a given lane of the PMA service interface below the RS-FEC sublayer, to compensate for skew between FEC lanes, and to identify RS-FEC codeword boundaries.

The alignment marker mapping function operates on a group of 20 aligned and reordered alignment markers. Let $\text{am_rx_x}<65:0>$ be the alignment marker for PCS lane x , $x = 0$ to 19, where bit 0 is the first bit transmitted. The alignment markers shall be mapped to $\text{am_rxmapped}<1284:0>$ in a manner that yields the same result as the following process.

For $x = 0$ to 19, $\text{amp_rx_x}<63:0>$ is constructed as follows:

- a) Set $y = 0$ when $x \leq 3$, set $y = 16$ when $x \geq 16$, otherwise set $y = x$.
- b) $\text{amp_rx_x}<23:0>$ is set to M_0 , M_1 , and M_2 as shown in Figure 82-9 (bits 25 to 2) using the values in Table 82-2 for PCS lane number y .
- c) $\text{amp_rx_x}<31:24> = \text{am_rx_x}<33:26>$.

- d) $\text{amp_rx_x}\langle 55:32 \rangle$ is set to M_4 , M_5 , and M_6 as shown in Figure 82–9 (bits 57 to 34) using the values in Table 82–2 for PCS lane number y .
- e) $\text{amp_rx_x}\langle 63:56 \rangle = \text{amp_rx_x}\langle 65:58 \rangle$.

This process replaces the fixed bytes of the alignment markers received, possibly with errors, with the values from Table 82–2. In addition it substitutes the fixed bytes of the alignment markers corresponding to PCS lanes 1, 2, and 3 with the fixed bytes for the alignment marker corresponding to PCS lane 0. Similarly, it substitutes the fixed bytes of the alignment markers corresponding to PCS lanes 17, 18, and 19 with the fixed bytes for the alignment marker corresponding to PCS lane 16. The variable bytes BIP or CD are unchanged. This process simplifies receiver synchronization since the receiver only needs to search for the fixed bytes corresponding to PCS lane 0 on each FEC lane.

Construct a matrix of 4 rows and 320 columns, am_rxpayloads , as shown in Figure 152–4. Given $i = 0$ to 3, $j = 0$ to 4, and $x = i + 4j$, the matrix is derived per the following expression:

$$\text{am_rxpayloads}\langle i, (64j + 63):64j \rangle = \text{amp_rx_x}\langle 63:0 \rangle$$

Given $i = 0$ to 3, $k = 0$ to 31, and $y = i + 4k$, am_rxmapped may then be derived from am_rxpayloads per the following expression:

$$\text{am_rxmapped}\langle (10y + 9):10y \rangle = \text{am_rxpayloads}\langle i, (10k + 9):10k \rangle$$

A 5-bit pad is appended to the mapped alignment markers to yield the equivalent of five 257-bit blocks. The pad bits, $\text{am_rxmapped}\langle 1284:1280 \rangle$, shall be set to the binary values 00101 and 11010 (the leftmost bit is assigned to the highest bit index) in an alternating pattern. In other words, if a pad value of 00101 is used for the current iteration of the mapping function, a value of 11010 is used in the next iteration and vice versa.

The result of the alignment marker mapping function is a deterministic mapping between alignment marker payloads and FEC lanes. The alignment marker payloads corresponding to PCS lanes 0, 4, 8, 12, and 16 are transmitted on FEC lane 0, the alignment marker payloads corresponding to PCS lanes 0, 5, 9, 13, and 16 are transmitted on FEC lane 1, and so on (see Figure 152–4).

As a result of this process, the BIP_3 and BIP_7 fields from normal alignment markers are carried across the link protected by FEC. These fields cannot be used to monitor errors on the link protected by FEC as 64B/66B to 256B/257B transcoding and Reed-Solomon encoding alters the bit sequence. However, these fields may again be used to monitor errors after the original bit sequence is restored, that is following Reed-Solomon decoding and 256B/257B to 64B/66B transcoding.

One group of aligned and reordered alignment markers are mapped every $20 \times 16 \times 384$ 66-bit blocks. This corresponds to 4096 Reed-Solomon codewords (refer to 152.5.3.7). The mapped alignment markers, $\text{am_txmapped}\langle 1284:0 \rangle$ shall be inserted as the first 1285 message bits to be transmitted from every 4096th codeword.

The first 257-bit block inserted after am_txmapped shall correspond to the four 66-bit blocks received on PCS lanes 0, 1, 2, and 3 that immediately followed the alignment marker on each respective lane.

FEC lane, <i>i</i>	Reed-Solomon symbol index, <i>k</i> (10-bit symbols)																																
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32
0	0	amp_rx_0				63	0	amp_rx_4				63	0	amp_rx_8				63	0	amp_rx_12				63	0	amp_rx_16				63	5-bit pad		
1	0	amp_rx_1				63	0	amp_rx_5				63	0	amp_rx_9				63	0	amp_rx_13				63	0	amp_rx_17				63	5-bit pad		
2	0	amp_rx_2				63	0	amp_rx_6				63	0	amp_rx_10				63	0	amp_rx_14				63	0	amp_rx_18				63	5-bit pad		
3	0	amp_rx_3				63	0	amp_rx_7				63	0	amp_rx_11				63	0	amp_rx_15				63	0	amp_rx_19				63	5-bit pad		

Legend:  = 5-bit pad

rx_scrambled

Figure 152–4—Alignment marker mapping to FEC lanes

152.5.3.7 Reed-Solomon encoder

The Reed-Solomon encoder implements the RS(544,514) FEC encoder described in 91.5.2.7. Since the encoder is used in the receive direction of transmission, the message symbols come from rx_scrambled and rx_ammapped rather than tx_scrambled and tx_ammapped.

152.5.3.8 Symbol distribution

Once the data has been Reed-Solomon encoded, it shall be distributed to 4 FEC lanes, one 10-bit symbol at a time in a round-robin distribution from the lowest to the highest numbered FEC lane. The distribution process is shown in Figure 152–5.

152.5.3.9 Receive bit ordering

The receive bit ordering is illustrated in Figure 152–5.

IECNORM.COM :: Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd 13:2022

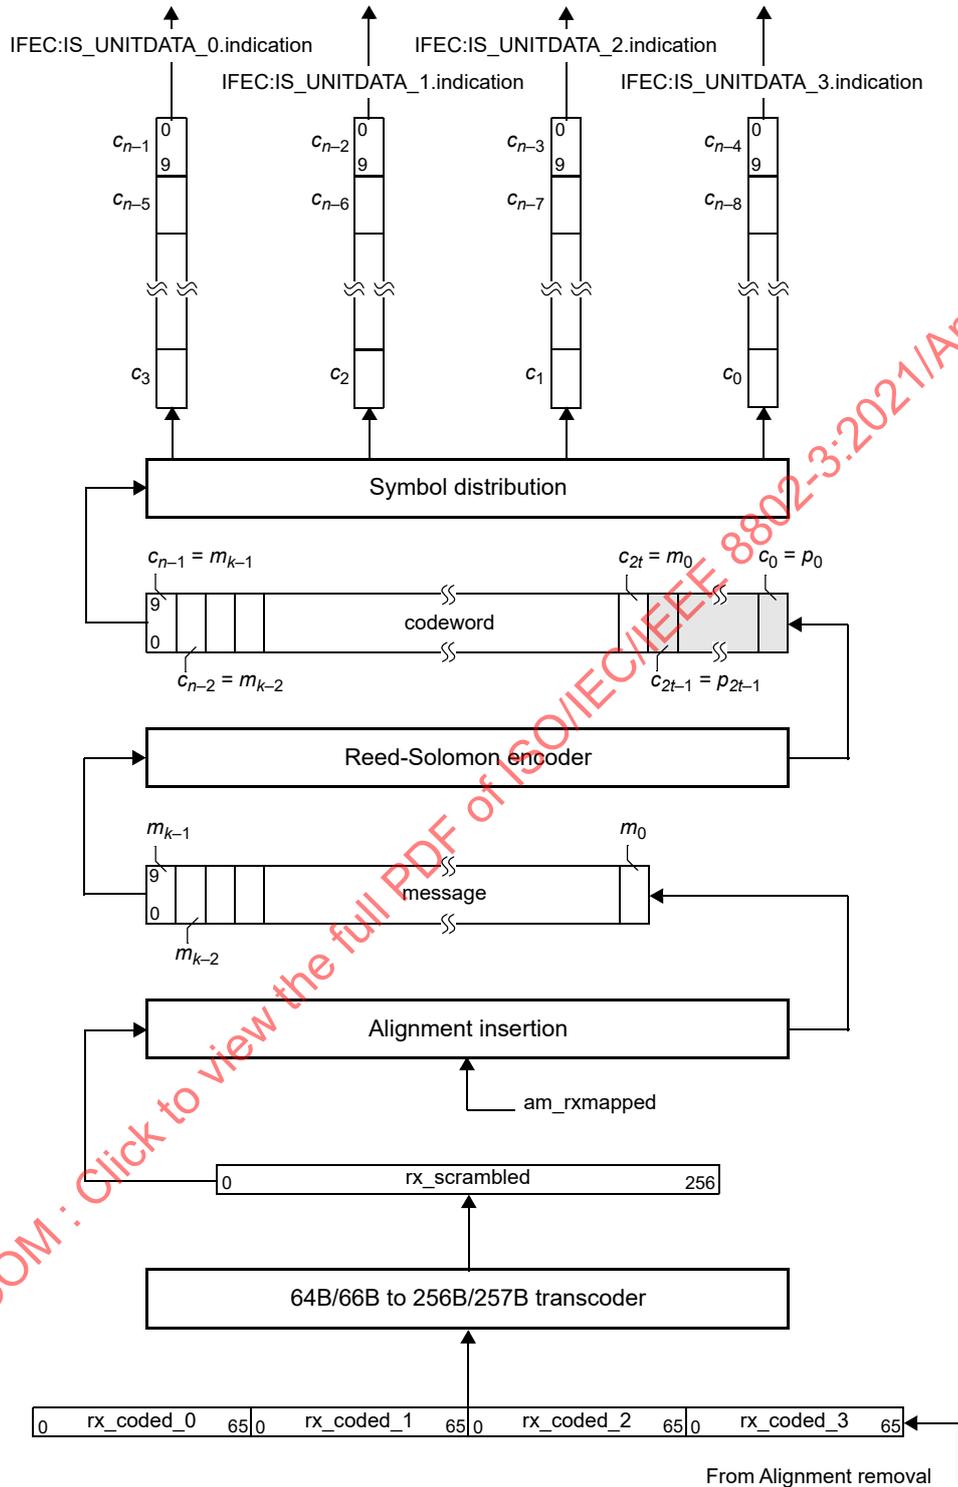


Figure 152-5—Receive bit ordering

152.5.4 Detailed functions and state diagrams

Most functions and state diagrams are identical to those described in 91.5.4, with only the modifications to variable naming that arise from reversing the Tx and Rx directions, and omission of EEE deep sleep capability due to the fact that this sublayer only supports optical PHY types.

152.5.4.1 State diagram conventions

The body of this subclause is composed of state diagrams, including the associated definitions of variables, functions, and counters. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

152.5.4.2 State variables

152.5.4.2.1 Variables

`all_locked`

A Boolean variable that is set to true when `amps_lock<x>` is true for all x and is set to false when `amps_lock<x>` is false for any x .

`amp_counter_done`

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

`amp_match`

Boolean variable that holds the output of the function `AMP_COMPARE`.

`amp_valid`

Boolean variable that is set to true if the received 64-bit block is a valid alignment marker payload. The alignment marker payload, mapped to an FEC lane according to the process described in 152.5.3.6, consists of 48 known bits and 16 variable bits (the BIP_3 or CD_3 field and its complement BIP_7 or CD_7 , see 82.2.7). The bits of the candidate block that are in the positions of the known bits in the alignment marker payload are compared on a nibble-wise basis (12 comparisons). If no more than 3 nibbles in the candidate block fail to match the corresponding known nibbles in the alignment marker payload, the candidate block is considered a valid alignment marker payload. Each FEC lane compares the candidate block to the alignment marker payload for PCS lane 0.

`amps_lock<x>`

Boolean variable that is set to true when the receiver has detected the location of the alignment marker payload sequence for a given lane on the IFEC service interface, where $x = 0:3$.

`current_pcs_l`

A variable that holds the PCS lane number corresponding to the current alignment marker payload that is recognized on a given lane of the FEC service interface. It is compared to the variable `first_pcs_l` to confirm that the location of the alignment marker payload sequence has been detected.

`cw_bad`

A Boolean variable that is set to true if the Reed-Solomon decoder (see 152.5.2.3) fails to correct the current FEC codeword and is set to false otherwise.

`deskew_done`

A Boolean variable that is set to true when `fec_enable_deskew` is set to true and the deskew process is completed. Otherwise, this variable is set to false.

`fec_align_status`

A Boolean variable set by the FEC alignment process to reflect the status of FEC lane-to-lane alignment. Set to true when all lanes are synchronized and aligned and set to false when the deskew process is not complete.

fec_alignment_valid

Boolean variable that is set to true if all FEC lanes are aligned. FEC lanes are considered to be aligned when amps_lock<x> is true for all x, each FEC lane is locked to a unique alignment marker payload sequence (see 152.5.3.6), and the FEC lanes are deskewed. Otherwise, this variable is set to false.

fec_enable_deskew

A Boolean variable that enables and disables the deskew process. Received bits may be discarded whenever deskew is enabled. It is set to true when deskew is enabled and set to false when deskew is disabled.

fec_lane

A variable that holds the FEC lane number (0 to 3) received on lane x of the IFEC service interface when amps_lock<x> = true. The FEC lane number is determined by the alignment marker payloads in the 2nd, 3rd, or 4th positions of the sequence based on the mapping defined in 152.5.3.6. The 48 bits that are in the positions of the known bits in the received alignment marker payload are compared to the expected values for a given payload position and FEC lane on a nibble-wise basis (12 comparisons). If no more than 3 nibbles in the candidate block fail to match the corresponding known nibbles for any payload position on a given FEC lane, then the FEC lane number is assigned accordingly.

fec_optional_states

Boolean variable that is always set to true to indicate that the optional states in the FEC synchronization state diagram in Figure 91–8 are implemented.

first_pcs

A variable that holds the PCS lane number that corresponds to the first alignment marker payload that is recognized on a given lane of the FEC service interface. It is compared to the PCS lane number corresponding to the second alignment marker payload that is tested.

reset

Boolean variable that controls the resetting of the Inverse RS-FEC sublayer. It is true whenever a reset is necessary including when reset is initiated from the MDIO, during power on.

restart_lock

Boolean variable that is set by the FEC alignment process to reset the synchronization process on all FEC lanes. It is set to true after 3 consecutive uncorrected codewords are received (3_BAD state) and set to false upon entry into the LOSS_OF_ALIGNMENT state. It is also set to true when 5 Alignment Markers in a row fail to match (5_BAD state).

rx_align_status

Boolean variable that is set by the alignment lock and deskew function (see 152.5.3.2).

signal_ok

Boolean variable that is set based on the most recently received value of FEC:IS_SIGNAL.indication(SIGNAL_OK). It is true if the value was OK and false if the value was FAIL.

slip_done

Boolean variable that is set to true when the SLIP requested by the synchronization state diagram has been completed indicating that the next candidate 64-bit block position can be tested.

test_amp

Boolean variable that is set to true when a candidate block position is available for testing and false when the FIND_1ST state is entered.

test_cw

Boolean variable that is set to true when a new FEC codeword is available for decoding and is set to false when the TEST_CW state is entered.

152.5.4.2.2 Functions

The AMP_COMPARE and SLIP functions as defined in 91.5.4.2.2 are used by the Inverse RS-FEC sublayer.

152.5.4.2.3 Counters

amp_bad_count

Counts the number of consecutive alignment markers that don't match the expected values for a given FEC lane.

amp_counter

This counter counts the 4096 FEC codewords that separate the ends of two consecutive normal alignment marker payload sequences. An FEC codeword is 1360 bits per FEC lane.

cw_bad_count

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

152.5.4.3 State diagrams

The FEC shall implement four synchronization processes as shown in Figure 91–8. The synchronization process operates independently on each lane. The synchronization state diagram determines when the FEC has detected the location of the alignment marker payload sequence in the received bit stream for a given lane of the service interface.

The FEC shall implement the alignment process as shown in Figure 91–9.

152.6 Inverse RS-FEC MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the Inverse RS-FEC. If MDIO is implemented, it shall map MDIO control bits to Inverse RS-FEC control variables as shown in Table 152–1, and MDIO status bits to Inverse RS-FEC status variables as shown in Table 152–2, and if a separated PMA (see 45.2.1) is connected to the FEC service interface it shall map additional MDIO status bits to additional RS-FEC status variables as shown in Table 152–3.

Table 152–1—MDIO/inverse RS-FEC control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	IFEC variable
IFEC bypass correction enable	IFEC control	1.2200.0	IFEC_bypass_correction_enable
IFEC bypass indication enable	IFEC control	1.2200.1	IFEC_bypass_indication_enable

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

152.6.1 IFEC_bypass_correction_enable

When this variable is set to one, the Reed-Solomon decoder performs error detection without error correction (see 152.5.2.3). When this variable is set to zero, the decoder also performs error correction. The default value of the variable is zero. This variable is mapped to the bit defined in 45.2.1.186aa.2 (1.2200.0).

Table 152–2—MDIO/Inverse RS-FEC status variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
IFEC bypass correction ability	IFEC status	1.2201.0	IFEC_bypass_correction_ability
IFEC bypass indication ability	IFEC status	1.2201.1	IFEC_bypass_indication_ability
IFEC high SER	IFEC status	1.2201.2	hi_ser
IFEC AM lock x , $x = 0$ to 3	IFEC status	1.2201.8:11	amps_lock< x >
IFEC align status	IFEC status	1.2201.14	IFEC_align_status
IFEC corrected codewords	IFEC corrected codewords counter	1.2202, 1.2203	IFEC_corrected_cw_counter
IFEC uncorrected codewords	IFEC uncorrected codewords counter	1.2204, 1.2205	IFEC_uncorrected_cw_counter
RS-FEC lane x mapping	IFEC lane mapping	1.2206	IFEC_lane_mapping< x >
IFEC symbol errors, lane 0 to 3	IFEC symbol error counter, lane 0 to 3	1.2210 to 1.2217	IFEC_symbol_error_counter_ i

Table 152–3—MDIO/Inverse RS-FEC status variable mapping for separated PMA

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
PCS align status	RS-FEC status	1.201.15	align_status
BIP error counter, lane 0 to 19	RS-FEC BIP error counter, lane 0 to 19	1.230 to 1.249	BIP_error_counter_ i
Lane x mapping	RS-FEC PCS lane x mapping	1.250 to 1.269	lane_mapping< x >
Block lock x	RS-FEC PCS alignment status 1 and 2	1.280, 1.281	block_lock< x >
Lane x aligned	RS-FEC PCS alignment status 3 and 4	1.282, 1.283	am_lock< x >

152.6.2 IFEC_bypass_indication_enable

This variable is set to one to bypass the error indication function (see 152.5.2.3) when this ability is supported. When this variable is set to zero, the decoder indicates errors to the PCS sublayer. This variable has no effect (the decoder does not bypass error indication) if IFEC_bypass_correction_enable (1.2200.0) is set to one. The default value of this variable is zero. This variable is mapped to the bit defined in 45.2.1.186aa.1 (1.2200.1).

152.6.3 IFEC_bypass_correction_ability

The Reed-Solomon decoder may have the option to perform error detection without error correction (see 152.5.2.3) to reduce the delay contributed by the Inverse RS-FEC sublayer. This variable is set to one to indicate that the decoder has the ability to bypass error correction. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.186ab.9 (1.2201.0).

152.6.4 IFEC_bypass_indication_ability

The Reed-Solomon decoder may have the option to bypass the error indication function (see 152.5.2.3) to reduce the delay contributed by the Inverse RS-FEC sublayer. This variable is set to one to indicate that the decoder has the ability to bypass the error indication function. The variable is set to zero if this ability is not supported. This variable is mapped to the bit defined in 45.2.1.186ab.8 (1.2201.1).

152.6.5 hi_ser

This variable is defined when the FEC_bypass_indication_ability variable is set to one. When FEC_bypass_indication_enable is set to one, this bit is set to one if the number of RS-FEC symbol errors in a window of 8192 codewords exceeds the threshold (see 152.5.2.3) and is set to zero otherwise. This variable is mapped to the bit defined in 45.2.1.186ab.7 (1.2201.2).

152.6.6 amps_lock<x>

These variables are assigned by the FEC synchronization state diagram shown in Figure 91–8 (see 152.5.4.3). They are mapped to the bits defined in 45.2.1.186ab (1.2201.8 to 1.2201.11 for FEC lanes 0 to 3, respectively).

152.6.7 IFEC_align_status

This variable is assigned by the FEC alignment state diagram shown in Figure 91–9 (see 152.5.4.3). It is mapped to the bit defined in 45.2.1.186ab.2 (1.2201.14).

152.6.8 IFEC_corrected_cw_counter

A corrected FEC codeword is a codeword that contained errors and was corrected.

The IFEC_corrected_cw_counter is a 32-bit counter that counts once for each corrected FEC codeword processed when IFEC_align_status is true. This variable is mapped to the registers defined in 45.2.1.186ac (1.2202, 1.2203).

152.6.9 IFEC_uncorrected_cw_counter

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

The IFEC_uncorrected_cw_counter is a 32-bit counter that counts once for each uncorrected FEC codeword processed when IFEC_align_status is true. This variable is mapped to the registers defined in 45.2.1.186ad (1.2204, 1.2205).

152.6.10 IFEC_lane_mapping<x>

When the Inverse RS-FEC transmit function detects and locks to an alignment marker payload on the service interface lane x , the FEC lane number corresponding to the detected alignment marker payload is assigned to the variable IFEC_lane_mapping< x >. These variables are mapped to the register defined in 45.2.1.186ae (1.2206).

152.6.11 IFEC_symbol_error_counter_{*i*}

IFEC_symbol_error_counter_{*i*}, where $i = 0$ to 3, is a 32-bit counter that counts once for each 10-bit symbol corrected on FEC lane i when IFEC_align_status is true. These variables are mapped to the registers defined in 45.2.1.186af and 45.2.1.186ag (1.2210 to 1.2217).

152.6.12 align_status

This variable is assigned the value of rx_align_status as defined by the PCS deskew state diagram shown in Figure 82–14 (see 152.5.3.2). It is mapped to the bit defined in 45.2.1.111 (1.201.15).

152.6.13 BIP_error_counter_{*i*}

BIP_error_counter_{*i*}, where $i = 0$ to 19, is a 16-bit counter that holds the BIP error count for PCS lane i as calculated by the RS-FEC receive function (see 152.5.3.4). These variables are mapped to the registers defined in 45.2.1.117 and 45.2.1.118 (1.230 to 1.249).

152.6.14 lane_mapping<*x*>

When the RS-FEC transmit function detects and locks to an alignment marker on FEC service interface lane x , the PCS lane number corresponding to the detected alignment marker is assigned to the variable lane_mapping< x >. These variables are mapped to the registers defined in 45.2.1.119 and 45.2.1.120 (1.250 to 1.269).

152.6.15 block_lock<*x*>

These variables are assigned by the block lock state diagram shown in Figure 82–12 (see 152.5.3.1). They are mapped to the registers defined in 45.2.1.121 and 45.2.1.122 (1.280 and 1.281).

152.6.16 am_lock<*x*>

These variables are assigned by the alignment marker lock state diagram shown in Figure 82–13 (see 152.5.3.2). They are mapped to the registers defined in 45.2.1.123 and 45.2.1.124 (1.282 and 1.283).

152.7 Protocol implementation conformance statement (PICS) proforma for Clause 152, Inverse RS-FEC sublayer³

152.7.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 152, Inverse RS-FEC sublayer, 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](#).

152.7.2 Identification

152.7.2.1 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).	

152.7.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ct-2021, Clause 152, Inverse RS-FEC sublayer
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 802.3ct-2021.)	

Date of Statement	
-------------------	--

³Copyright release for PICS proformas: Users of this standard may freely reproduce the PICS proforma in this subclause so that it can be used for its intended purpose and may further publish the completed PICS.

152.7.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
DC	Delay constraints	152.4	Conforms to delay constraints specified in 152.4	M	Yes []
*MD	MDIO capability	152.6	Registers and interface supported	O	Yes [] No []
*BEC	Bypass error correction	152.5.2.3	Capability is supported	O	Yes [] No []
*BEI	Bypass error indication	152.5.2.3	Capability is supported	O	Yes [] No []

152.7.4 PICS proforma tables for Inverse RS-FEC sublayer

152.7.4.1 Transmit function

Item	Feature	Subclause	Value/Comment	Status	Support
TF1	Skew tolerance	152.5.2.1	Maximum Skew of 49 ns between FEC lanes and a maximum Skew Variation of 400 ps	M	Yes []
TF2	Lane reorder	152.5.2.2	Order the FEC lanes according to the FEC lane number	M	Yes []
TF3	Reed-Solomon decoder	152.5.2.3	RS(544,514) corrects any combination of up to $t = 15$ symbol errors in a codeword unless error correction bypassed	M	Yes []
TF4	Reed-Solomon decoder uncorrected codeword detection	152.5.2.3	Capable of indicating when a codeword was not corrected	M	Yes []
TF5	Error indication function	152.5.2.3	Corrupts 66-bit block synchronization headers for uncorrected errored codewords (or errored codewords when correction is bypassed)	M	Yes []
TF6	Error indication when error correction is bypassed	152.5.2.3	Error indication is not bypassed	BEI:M	Yes [] N/A []
TF7	Error monitoring while error indication is bypassed	152.5.2.3	When the number of symbol errors in a block of 8192 codewords exceeds 6380, corrupt 66-bit block synchronization headers	BEI:M	Yes [] N/A []
TF8	Alignment marker removal	152.5.2.4	am_txmapped removed prior to transcoding	M	Yes []
TF9	256B/257B to 64B/66B transcoder	152.5.2.5	tx_coded_j<65:0>, j = 0 to 3 constructed per 152.5.2.5	M	Yes []

Item	Feature	Subclause	Value/Comment	Status	Support
TF10	Block distribution	152.5.2.6	One 66-bit block at a time in a round-robin fashion from the lowest to the highest numbered PCS lane	M	Yes []
TF11	Alignment marker mapping	152.5.2.7	Map to am_tx_x, x = 0 to 19 per 152.5.2.7	M	Yes []
TF12	Alignment marker insertion point	152.5.2.7	Alignment markers immediately followed by the 66-bit blocks derived from the 256-bit blocks immediately following am_txmapped	M	Yes []

152.7.4.2 Receive Function

Item	Feature	Subclause	Value/Comment	Status	Support
RF1	Skew tolerance	152.5.3.2	Maximum Skew of 49 ns between PCS lanes and a maximum Skew Variation of 400 ps	M	Yes []
RF2	Lane reorder	152.5.3.3	Order the PCS lanes according to the PCS lane number	M	Yes []
RF3	64B/66 to 256B/257B transcoder	152.5.3.5	rx_coded<256:0> constructed per 152.5.3.5	M	Yes []
RF4	257-bit block transmission order	152.5.3.5	First bit transmitted in bit 0	M	Yes []
RF5	Alignment marker mapping	152.5.3.6	Map to am_rx-mapped<1284:0> per 152.5.3.6	M	Yes []
RF6	Pad value	152.5.3.6	Binary values 00101 and 11010 (the leftmost bit is assigned to the highest bit index) in an alternating pattern	M	Yes []
RF7	Alignment marker insertion	152.5.3.6	First 1285 message bits to be transmitted from every 4096th codeword	M	Yes []
RF8	Alignment marker insertion point	152.5.3.6	First 257-bit block inserted after am_rxmapped corresponds to the four 66-bit blocks received on PCS lanes 0, 1, 2, and 3 that immediately followed the alignment marker on each respective lane	M	Yes []
RF9	Reed-Solomon encoder	152.5.3.7	RS(544,514)	M	Yes []
RF10	Symbol distribution	152.5.3.8	Distributed to 4 FEC lanes, one 10-bit symbol at a time in a round-robin distribution from the lowest to the highest numbered FEC lane	M	Yes []

152.7.4.3 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	SLIP function	152.5.4.2.2	Ensures that all possible block positions are evaluated	M	Yes []
SD2	Synchronization process	152.5.4.3	One instance per FEC lane per Figure 91-8	M	Yes []
SD3	Alignment process	152.5.4.3	Per Figure 91-9	M	Yes []
SD4	FEC synchronization enhancement	152.5.4.3	Check AMs after synchronization is achieved per Figure 91-8	M	Yes []

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd.13:2022

153. SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs

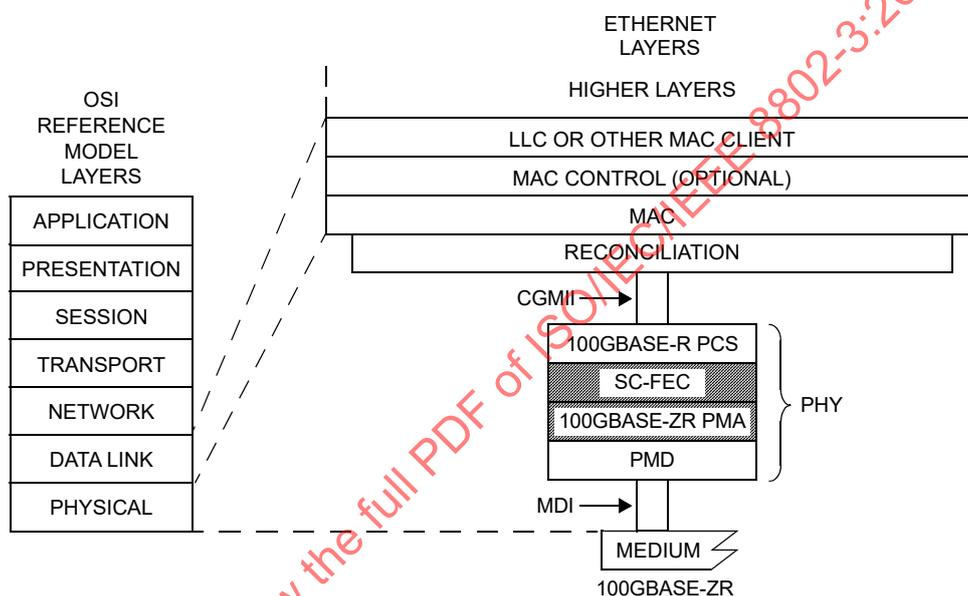
153.1 Overview

153.1.1 Scope

This clause specifies a staircase FEC (SC-FEC) and PMA sublayer for 100GBASE-ZR PHYs.

153.1.2 Position of SC-FEC and 100GBASE-ZR PMA in the 100GBASE-R sublayers

Figure 153–1 shows the relationship of the SC-FEC and 100GBASE-ZR PMA to the ISO/IEC Open System Interconnection (OSI) reference model.



CGMII = 100 Gb/s MEDIA INDEPENDENT INTERFACE
 LLC = LOGICAL LINK CONTROL
 MAC = MEDIA ACCESS CONTROL
 MDI = MEDIUM DEPENDENT INTERFACE
 PCS = PHYSICAL CODING SUBLAYER
 PHY = PHYSICAL LAYER DEVICE
 PMA = PHYSICAL MEDIUM ATTACHMENT
 PMD = PHYSICAL MEDIUM DEPENDENT
 SC-FEC = STAIRCASE FORWARD ERROR CORRECTION

Figure 153–1—SC-FEC and 100GBASE-ZR PMA relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and the IEEE 802.3 Ethernet model

153.2 SC-FEC sublayer

153.2.1 FEC service interface

This subclause specifies the services provided by the SC-FEC sublayer.

The FEC service interface is provided to allow the PCS, Inverse RS-FEC, or PMA to transfer information to and from the SC-FEC. The service interface is described in an abstract manner and does not imply any particular implementation. The FEC service interface is an instance of the inter-sublayer service interface defined in 80.3. The service interface primitives are summarized as follows:

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

The FEC:IS_UNITDATA_ *i* primitives are defined for $i = 0$ to 19. The PCS, Inverse RS-FEC, or PMA continuously sends 20 parallel bit streams to the SC-FEC sublayer, each at a nominal signaling rate of 5.15625 GBd. The SC-FEC sublayer continuously sends 20 parallel bit streams to the PCS, Inverse RS-FEC, or PMA, one per lane, each at a nominal signaling rate of 5.15625 GBd.

The SIGNAL_OK parameter of the FEC:IS_SIGNAL.indication primitive can take one of two values: OK or FAIL. The value is set to OK when the FEC receive function has identified codeword boundaries as indicated by fec_align_status equal to TRUE. The value is set to FAIL when the FEC receive function is unable to reliably establish codeword boundaries as indicated by fec_align_status equal to FALSE. When SIGNAL_OK is FAIL, the rx_bit parameters of the FEC:IS_UNITDATA_ *i*.indication primitives are undefined.

The PCS or Inverse RS-FEC may be connected to the SC-FEC using an optional instantiation of the PMA service interface (see Annex 83A, Annex 83B, Annex 83D, and Annex 83E) in which case a PMA (see Clause 83) or Inverse FEC (see Clause 152) is a client of the FEC service interface.

153.2.2 Delay constraints

The maximum delay contributed by the SC-FEC sublayer (sum of transmit and receive delays at one end of the link) shall be no more than 1 827 840 bit times (3570 pause_quanta or 18 278.4 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

153.2.3 Functions within the SC-FEC sublayer

153.2.3.1 Functional block diagram

A functional block diagram of the SC-FEC sublayer is shown in Figure 153–2.

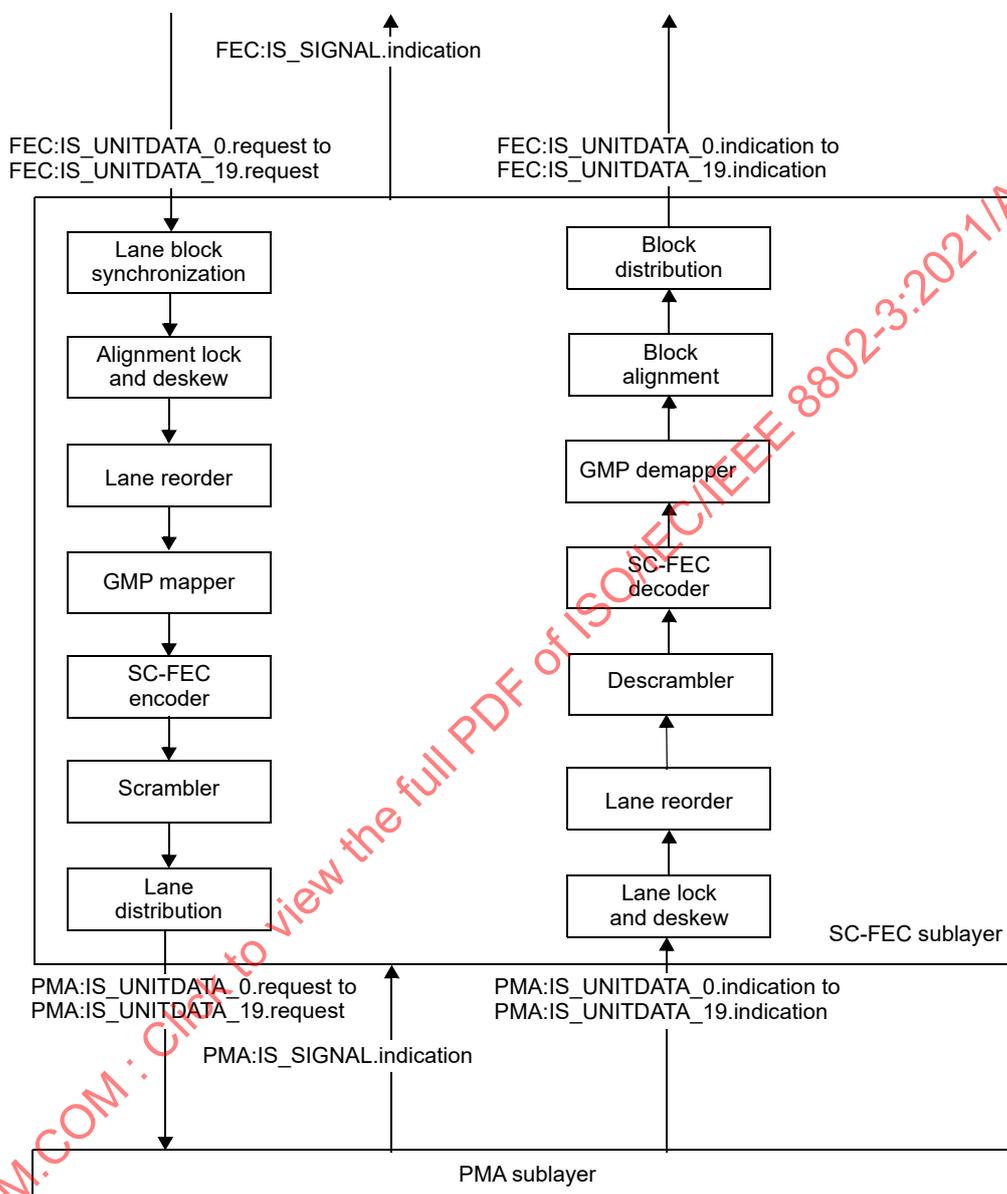


Figure 153–2—SC-FEC functional block diagram

153.2.3.2 Transmit function

153.2.3.2.1 Lane block synchronization

The SC-FEC transmit function forms 20 bit streams by concatenating the bits from each of the 20 `FEC:IS_UNITDATA_i.request` primitives in the order they are received. It obtains lock to the 66-bit blocks

in each bit stream using the sync headers and outputs 66-bit blocks. Block lock is obtained as specified in the block lock state diagram shown in Figure 82–12.

153.2.3.2.2 Alignment lock and deskew

Once the SC-FEC transmit function achieves block lock on a PCS lane, it then begins obtaining alignment marker lock as specified by the alignment marker lock state diagram shown in Figure 82–13. This process identifies the PCS lane number received on a particular lane of the service interface. After alignment marker lock is achieved on all 20 lanes, all inter-lane Skew is removed as specified by the PCS deskew state diagram shown in Figure 82–14. The SC-FEC transmit function shall support a maximum Skew of 49 ns between PCS lanes and a maximum Skew Variation of 400 ps. Skew and Skew Variation are defined in 80.5.

153.2.3.2.3 Lane reorder

PCS lanes can be received on different lanes of the service interface from which they were originally transmitted due to Skew between lanes and multiplexing by the PMA. The SC-FEC transmit function shall order the PCS lanes according to the PCS lane number. The result of this process is a logically serial stream of 66B block encoded information, including alignment markers, at a nominal signaling rate of 103.125 Gb/s.

153.2.3.2.4 GMP mapper

The generic mapping procedure (GMP) mapper inserts the serialized and deskewed stream of 66B blocks into an SC-FEC frame. The frame is similar to the OTU4 frame specified in ITU-T G.709 with the insertion of an SC-FEC parity specified in ITU-T G.709.2 Annex A. The frame consists of 16 320 octets, which for the purpose of description are illustrated as a four row, 4080 column structure with a logical transmission order left to right, top to bottom. This frame is illustrated in Figure 153–3.

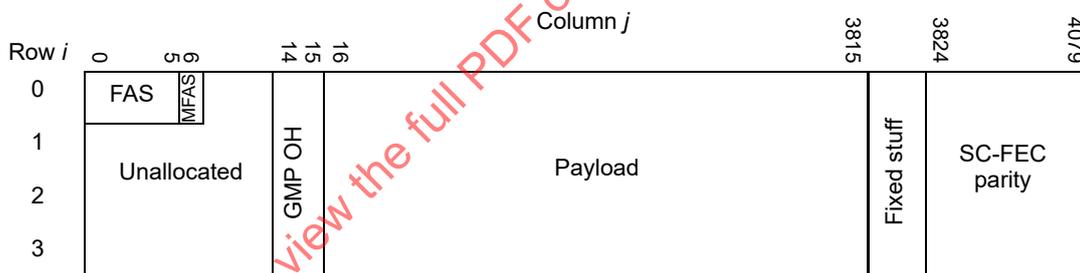


Figure 153–3—SC-FEC frame

The majority of the frame is scrambled prior to transmission (see 153.2.3.2.6). The information in the SC-FEC frame includes the following, described before scrambling:

- 1) The frame alignment signal (FAS) is similar in concept to the codeword marker described in Clause 108, being a fixed bit pattern occurring in the data stream at regular intervals allowing the receiver to locate and distinguish various elements in the transmitted data stream such as the payload and the FEC parity. The FAS is the following fixed bit pattern, transmitted left to right:
 1111 0110 1111 0110 1111 0110 0010 1000 0010 1000 0010 1000
- 2) The multi-frame alignment signal (MFAS) is a field that counts from 0 to 255, encoded with the most significant bit transmitted first. This is used for aligning SC-FEC base blocks with the SC-FEC frame and synchronizing the SC-FEC error decorrelator.
- 3) The unallocated portion of the frame corresponds to octet positions that are used for OTN path and section overhead in ITU-T G.709.2. These byte positions are not used for 100GBASE-ZR. They are transmitted with the value zero and are ignored on receipt.

- 4) The GMP mapping overhead (GMP OH) is encoded in the eight octets from columns 14 and 15 of each of the four rows of the SC-FEC frame.
- 5) The 100GBASE-R PCS payload is GMP mapped in the area of the SC-FEC frame in the octets in columns 16 through 3815 in each of the four rows as shown in Figure 153–3.
- 6) The octets in the fixed stuff area of the frame are transmitted with the value zero and are ignored on receipt.
- 7) The SC-FEC parity area of the frame is allocated for FEC parity to be inserted as described in 153.2.3.2.5.

GMP is a generic mechanism that uses a sigma/delta distribution algorithm to accommodate an arbitrary signaling rate difference between a payload and the space in which it is carried. The principles of the GMP mapper are described in ITU-T G.709 (06/2020) Annex D, with details of the encoding of the GMP overhead in ITU-T G.709 Clause 19.4.3.2.

The payload area of the SC-FEC frame is divided into 190 GMP words of 80 octets each. Each 80-octet GMP word is either filled with data (the logically serialized 66B encoded stream produced according to 153.2.3.2.3) or stuff, which is transmitted as zero prior to scrambling and ignored on receipt.

The 66B encoded data is a logically serial stream at a rate of $103.125 \text{ Gb/s} \pm 100 \text{ ppm}$. The payload area of the SC-FEC frame has a capacity of $(255/227) \times (3800/4080) \times 99.5328 \text{ Gb/s} \pm 20 \text{ ppm}$ ($\sim 104.1367 \text{ Gb/s}$). The clocks for the PCS and the SC-FEC frame are independent. This results in an average number of 80-octet GMP words filled per SC-FEC frame between ~ 188.1315 and ~ 188.1766 . Since the number of GMP words filled in a frame is always an integer, a given SC-FEC frame will have either 188 or 189 filled GMP words. While the GMP mechanism is generic, the particular clock rates and tolerances for this application result in a small number of cases, allowing the positions of data and stuff to be pre-computed. For either 188 or 189 GMP words filled in the SC-FEC frame, the GMP word consisting of the octets from column 16 to column 95 of row 0 will always be stuff (filled with zero rather than 66B encoded data prior to scrambling). In the case of 188 GMP words filled in the SC-FEC frame, the GMP word consisting of the octets from column 16 to column 95 of row 2 will also be stuff. Given the range of clock rate differences that may exist between the 66B encoded data and the SC-FEC frame payload capacity, there will tend to be between 4 and 7 frames with 188 filled GMP words between every frame with 189 filled GMP words. There is not a circumstance where there would be two consecutive frames containing 189 filled GMP words.

The GMP overhead is illustrated in Figure 153–4.

Row	Column 14								Column 15							
	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
0	0	0	0	D9	D8	D7	D6	D5	C13	C12	C11	C10	C9	C8	C7	C6
1	0	0	0	D4	D3	D2	D1	D0	C5	C4	C3	C2	C1	C0	II	DI
2	0	0	0	CRC-5				CRC-8								
3	0x00								0x00							

Figure 153–4—GMP overhead encoding

The information in rows 0 and 1 of column 15 indicates how many (188 or 189) GMP words are filled in the next SC-FEC frame. While GMP is a generic mechanism, the particular use on this interface leaves only a small number of ways to fill in this information. Since the number of GMP words to be filled in a frame is normally expected to vary within a small range, special coding is used to indicate when the number of filled GMP words is the same, one greater, or one fewer than in the previous frame. The following cases exist for this interface:

- The number of GMP words to be filled in the next SC-FEC frame is the same as the number of GMP words filled in this frame. This can only occur in the case of 188 filled GMP words, as the ratio of the two clock rates does not provide a case where two consecutive frames will contain 189 filled GMP words. The number 188 is encoded in bits C0 through C13, with C0 being the least significant bit and C13 being the most significant bit. The increment indicator (II bit) and decrement indicator (DI bit) are set to zero.
- The number of GMP words to be filled in the next SC-FEC frame is one greater than the number of GMP words filled in this SC-FEC frame. This will only occur in the case where the number of GMP words filled in this SC-FEC frame is 188 and the number of GMP words to be filled in the next SC-FEC frame is 189. This is signaled by inverting all of the odd-numbered C bits (C13, C11, C9, C7, C5, C3, and C1) from the value in the previous frame, setting the increment indicator (II bit) to one, and setting the decrement indicator (DI bit) to zero.
- The number of GMP words to be filled in the next SC-FEC frame is one fewer than the number of GMP words filled in this SC-FEC frame. This will only occur in the case where the number of GMP words filled in this SC-FEC frame is 189 and the number of GMP words to be filled in the next SC-FEC frame is 188. This is signaled by inverting all of the even-numbered C bits (C12, C10, C8, C6, C4, C2, C0) from the numeric value of GMP words filled in the previous frame (189), setting the decrement indicator (DI bit) to one, and setting the increment indicator (II bit) to zero.

The bit values for these three cases are shown in Table 153–1.

Table 153–1—Encoding of GMP words in next SC-FEC frame

This frame	Next frame	C13	C12	C11	C10	C9	C8	C7	C6	C5	C4	C3	C2	C1	C0	II	DI
188	188	0	0	0	0	0	0	1	0	1	1	1	1	0	0	0	0
	189	1	0	1	0	1	0	0	0	0	1	0	1	1	0	1	0
189	188	0	1	0	1	0	1	1	1	1	0	1	0	0	0	0	1

While there is an 8-bit Hamming distance between valid values that can go in rows 0 and 1 of column 15, this information is also protected by a CRC-8 inserted into row 2 of column 15. The generator polynomial for this CRC-8 is $x^8 + x^3 + x^2 + 1$, and is calculated as follows:

- The octets from rows 0 and 1 of column 15 are taken in network octet order, most significant/first transmitted bit first, to form a 16-bit pattern representing the coefficients of a polynomial $M(x)$ of degree 15.
- $M(x)$ is multiplied by x^8 and divided (modulo 2) by the generating polynomial, producing a remainder $R(x)$ of degree 7 or less.
- The coefficients of $R(x)$ are considered to be an 8-bit sequence, where x^7 is the most significant bit. This 8-bit sequence is the CRC-8 where the first bit of the CRC-8 to be transmitted is the coefficient of x^7 and the last bit transmitted is the coefficient of x^0 .

To reduce the mapping jitter, extra timing information is provided to allow the Tx to inform the Rx how many octets of 66B encoded PCS data are received in each SC-FEC frame period. This can be understood as the number of extra octets (or portions of octets) that are held in the Tx buffer, whose fill level varies with the clock offset of the incoming 66B blocks and depending on whether 188 or 189 GMP words are filled in a given SC-FEC frame. As an example, consider a case where both the PCS and SC-FEC clocks are running at approximately nominal rate, and an average of 15 052 octets of 66B encoded PCS data are to be transmitted in each SC-FEC frame period. This is 12 octets more than 188 filled GMP words of 80 octets.

Assume in the first frame period, where 188 GMP words are filled, that there are 15 excess octets in the Tx buffer, fewer than required to fill 189 GMP words. 188×80 octets of 66B encoded data are transmitted, and in the next SC-FEC frame period the Tx buffer fill has grown to 27 excess octets. The Tx buffer fill grows by 12 octets each frame period when 188 GMP words are filled, so grows from 39 to 51 to 63 to 75 octets. Then, in the seventh SC-FEC frame, there are sufficient extra octets in the Tx buffer to fill 189×80 octets in the SC-FEC frame, so 189×80 octets are transmitted, and the buffer fill level drops to $75 + 12 - 80 = 7$ octets.

This extra timing information, essentially the number of excess octets in the Tx buffer, is signaled from the Tx to the Rx in rows 0 and 1 of column 14 of the SC-FEC frame. This is signaled as a 10-bit value in bits D0 through D9, with D0 being the least significant bit. This field may take values from 0 to 80 to reflect the number of excess full or partial octets in the Tx buffer. This information allows the Rx to know how many octets to distribute from its buffer in each SC-FEC frame period, and hence allows maintaining one-octet timing granularity even though the mapping granularity into the SC-FEC frame is 80 octet groups.

The bits D0 through D9 are protected by a CRC-5, which is calculated using a generator polynomial of $x^5 + x + 1$. This is calculated by the following steps:

- The bits D9 through D0 are taken in transmission order, most significant bit first, to form a 10-bit pattern representing the coefficients of a polynomial $M(x)$ of degree 9.
- $M(x)$ is multiplied by x^5 and divided (modulo 2) by the generator polynomial, producing a remainder $R(x)$ of degree 4 or less.
- The coefficients of $R(x)$ are considered to be a 5-bit sequence, where x^4 is the most significant bit.
- This 5-bit sequence is the CRC-5 where the first bit of the CRC-5 to be transmitted is the coefficient of x^4 and the last bit transmitted is the coefficient of x^0 .

The third row of the GMP overhead is transmitted as zero and ignored on receipt.

153.2.3.2.5 SC-FEC encoder

The SC-FEC parity is calculated and inserted into the octets in columns 3824 through 4079 of each row of the FEC frame shown in Figure 153–3. The SC-FEC code is a blockwise recursively encoded 512×510 staircase code, sandwiched between a $30\,592 + 2048$ bit-wide optimized error decorrelator interleaver and error correlator de-interleaver. Computation of the FEC parity is described in ITU-T G.709.2 Annex A. The 512×510 code block distributes over two of the FEC frames shown in Figure 153–3.

153.2.3.2.6 Scrambler

The operation of the scrambler shall be functionally identical to a frame synchronous scrambler of length 65 535 operating at the bit-rate of the SC-FEC frame. The generating polynomial shall be $1 + x + x^3 + x^{12} + x^{16}$. The scrambler shall be reset to 0xFFFF at the first transmitted bit of the MFAS octet. The FAS octets are not scrambled. Figure 153–5 provides a functional block diagram of the frame synchronous scrambler.

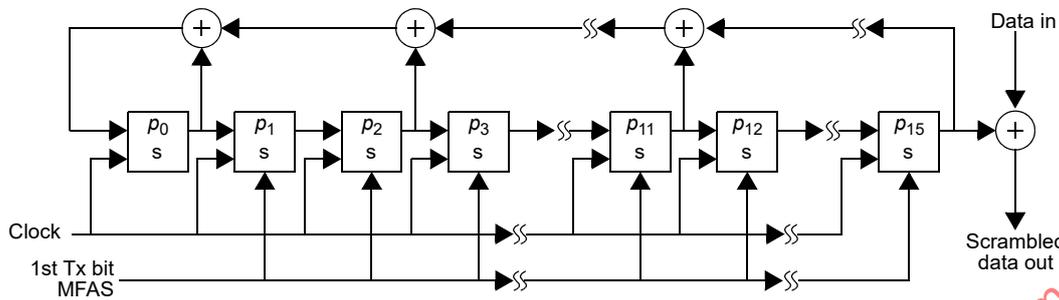


Figure 153-5—Frame synchronous scrambler

153.2.3.2.7 Lane distribution

The octets of the FEC frame are distributed to 20 FEC lanes as follows: The 6th octet of the FAS is replaced by a lane counter that increments from 0 to 239 on successive frames, with the most significant bit of the counter transmitted first. The lane counter 0 position shall be aligned with MFAS = 0 position every 3840 (the least common multiple of 240 and 256) frame periods. The counter resets to 0 on the first frame following the frame where the counter reaches the value 239.

Starting from the beginning of a FEC frame where the lane counter modulo 20 equals zero, the FEC frame is distributed to 20 FEC lanes 16 octets at a time in a round-robin manner as illustrated in Figure 153-6.

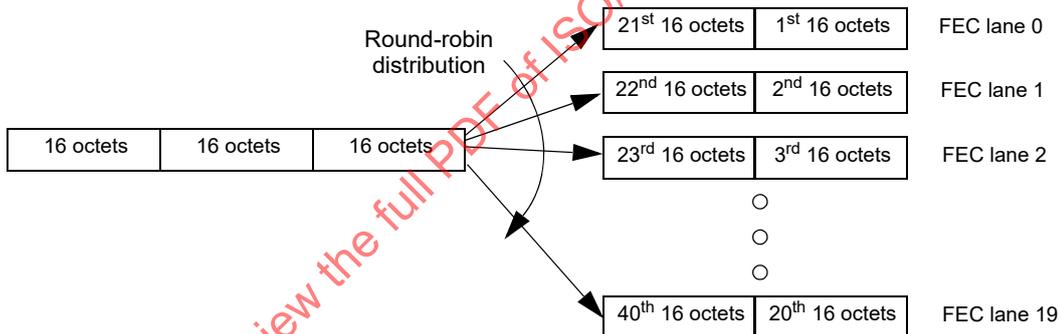


Figure 153-6—Lane distribution

51 groups of 16 octets are distributed from the FEC frame (consisting of 4080×4 octets) to each of the 20 FEC lanes.

At each FEC frame boundary, the assignment of 16-octet groups to FEC lanes is rotated, such that for each FEC frame where the lane counter modulo 20 equals 1, the first 16 octets of the FEC frame is distributed to FEC lane 1, the second 16 octets of the FEC frame is distributed to FEC lane 2, and so forth, with the 20th 16 octets of the FEC frame distributed to FEC lane 0. This pattern continues, with the final rotation occurring where the lane counter modulo 20 equals 19, where the first 16 octets of the FEC frame is distributed to FEC lane 19.

NOTE—This distribution to FEC lanes follows the distribution method of OTU4 to parallel lanes described in ITU-T G.709 Annex C.

153.2.3.3 Receive function

153.2.3.3.1 Lane lock and deskew

The SC-FEC receive function forms 20 bit streams by concatenating the bits from each of the 20 PMA:IS_UNITDATA_*i*.indication primitives in the order they are received. It obtains lock to the FAS as specified by the SC-FEC synchronization state diagram shown in Figure 153–7. After frame alignment lock is achieved on all 20 lanes, all inter-lane Skew is removed as specified by the SC-FEC deskew state diagram shown in Figure 153–8. The FEC receive function shall support a maximum Skew of 180 ns between FEC lanes and a maximum Skew Variation of 4 ns.

NOTE—An alternative method for frame alignment of OTU4 distributed to parallel lanes described in ITU-T G.709 Annex C is described in ITU-T G.798 [B48]. Equipment using this alternative method of frame alignment can interconnect with equipment implementing the SC-FEC synchronization state diagram shown in Figure 153–7, with each end of the link identifying and deskewing lanes according to its own method. The methods should not be combined at the same end of the link: in particular, if the ITU-T G.798 [B48] process is used, SIGNAL_OK=FALSE should not be indicated when in the out-of-frame (OOF) or out-of-recovery (OOR) states.

153.2.3.3.2 Lane reorder

Transmit FEC lanes can be received on different lanes of the service interface from which they were originally transmitted due to Skew between lanes and multiplexing by the PMA. The receive SC-FEC shall order the received FEC lanes according to the FEC lane number, which is the 6th octet of the FAS (inserted as per 153.2.3.2.7) modulo 20. As the reordered lanes are interleaved in the original order, the FEC lane number is replaced by the original final octet of the FAS, which has a value of 0010 1000.

153.2.3.3.3 De-scrambler

The de-scrambler reverses the effects of the scrambler described in 153.2.3.2.6. Since the scrambler is frame synchronous, this is accomplished by XORing each bit after the FAS, which is not scrambled, with the same sequence resulting from the generating polynomial used by the scrambler.

153.2.3.3.4 SC-FEC decoder

Errors are corrected by applying the process described in ITU-T G.709.2 Annex A. This operates by passing the received information and parity bits through an error decorrelator interleaver process, a parity compute, compare, and error correction process, and an error decorrelator deinterleaver process. The FEC parity is removed from the frame shown in Figure 153–3 after error correction.

The SC-FEC decoder indicates errors to the PCS sublayer by replacing all 66-bit blocks (extracted from the SC-FEC frame using the process described in 153.2.3.3.5) that are fully or partially contained within an uncorrected codeword (including uncorrectable codewords, or potentially mis-corrected or not fully corrected codewords) with error control blocks. An error control block has control block type 0x1E and carries eight /E/ control characters. The marking includes a 66-bit block that begins in the SC-FEC codeword preceding the uncorrected SC-FEC codeword and ending in the uncorrected codeword, and a 66-bit block that begins in the uncorrected codeword and ends in the next codeword. The probability that the SC-FEC decoder fails to replace 66-bit blocks in an uncorrectable codeword is expected to be less than 10^{-6} .

153.2.3.3.5 GMP demapper

The GMP demapper extracts the deskewed and serialized stream of 66B blocks that was inserted according to the process described in 153.2.3.2.4 from the SC-FEC frame.

The principles of the GMP demapper are described in ITU-T G.709 Annex D, with details of the encoding of the GMP overhead in ITU-T G.709 Clause 19.4.3.2.

The GMP demapper extracts 188 or 189 groups of 80 octets of 66B encoded data from each SC-FEC frame (frame n), controlled by the GMP overhead (see Table 153–1) in the previous SC-FEC frame (frame $n - 1$). If a C13:C0 value other than 188 or 189, or DI=1 and II=1 is received, the GMP demapper behavior is undefined. The number of 80-octet groups demapped from SC-FEC frame n , together with the change in value of the extra timing information (see 153.2.3.2.4) in between SC-FEC frames $n - 1$ and $n - 2$ are used to calculate how many octets of 66B encoded data are sent towards the FEC service interface in each SC-FEC frame period.

153.2.3.3.6 Block alignment

Block lock is obtained on the bit sequence extracted by the GMP demapper using the block lock state diagram shown in Figure 82–12.

153.2.3.3.7 Block distribution

Once block alignment has been achieved per 153.2.3.3.6, 66B blocks are distributed round robin to the lanes of the service interface as described in 82.2.6.

153.2.4 Detailed functions and state diagrams

The body of this subclause includes state diagrams, including the associated definitions of variables, functions, and counters. Should there be a discrepancy between a state diagram and descriptive text, the state diagram prevails.

The notation used in the state diagrams follows the conventions of 21.5. The notation ++ after a counter or integer variable indicates that its value is to be incremented.

153.2.4.1 State variables

153.2.4.1.1 Variables

all_locked

A Boolean variable that is set to TRUE when fas_lock< x > is TRUE for all x and is set to FALSE when fas_lock< x > is FALSE for any x .

current_fecl

A variable that holds the FEC lane identification octet corresponding to the current FAS that is recognized on a given lane of the PMA service interface. The value of this variable modulo 20 represents the FEC lane number. It is compared to the expected value based on the variable prev_fecl to confirm that the location of the FAS has been detected. This value is interpreted with the most significant bit transmitted first.

fas_counter_done

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

fas_match

Boolean variable that holds the output of the function FAS_COMPARE.

fas_valid

Boolean variable that is set to TRUE if the received 6-octet sequence is a valid FAS. The FAS consists of five known octets and one variable octet. The sequence is considered to be valid if at least four of the first five octets match the known bits of the pattern described in 153.2.3.2.4, and the 6th octet represents a numerical value in the range 0 to 239 with the most-significant bit transmitted first.

fas_lock< x >

Boolean variable that is set to TRUE when the receiver has detected the location of the FAS for a given lane on the PMA service interface, where $x = 0:19$.

fec_align_status

A Boolean variable set by the FEC alignment process to reflect the status of FEC lane-to-lane alignment. Set to TRUE when all lanes are synchronized and aligned and set to FALSE when the deskew process is not complete.

fec_alignment_valid

Boolean variable that is set to TRUE if all FEC lanes are aligned. FEC lanes are considered to be aligned when `fas_lock<x>` is TRUE for all x , each FEC lane has a unique lane number, and the FEC lanes are deskewed. Otherwise, this variable is set to FALSE.

fec_enable_deskew

A Boolean variable that indicates the enabling and disabling of the deskew process. Data may be discarded whenever deskew is enabled. TRUE when deskew is enabled. FALSE when deskew is disabled.

fec_lane

A variable that holds the FEC lane number (0 to 19) received on lane x of the PMA service interface when `fas_lock<x>` = TRUE. The FEC lane number is determined by the 6th octet of the FAS, interpreted with the most significant bit transmitted first, modulo 20. It is set to `current_fecl` modulo 20 when the expected FAS is found in the expected location in a second consecutive SC-FEC frame.

FEC_lane_mapping<x>

This variable indicates which FEC lane is received on lane x of the PMA service interface when `fas_lock<x>` = TRUE, where $x = 0:19$.

prev_fecl

A variable that holds the value of the lane identification octet (or the expected value of the lane identification octet) from the previous SC_FEC frame. The value is interpreted with the most significant bit transmitted first. The lane represented is this value modulo 20. It is used to calculate the expected value of the lane identification octet in the next SC-FEC frame period that is tested.

reset

Boolean variable that controls the resetting of the SC-FEC sublayer. It is TRUE whenever a reset is necessary including when reset is initiated from the MDIO and during power on.

restart_lock

Boolean variable that is set by the FEC alignment process to reset the synchronization process on all FEC lanes. It is set to TRUE when 15 FASs in a row fail to match (15_BAD state).

rx_align_status

Boolean variable that is set by the lane lock and deskew function (see 153.2.3.3.1).

signal_ok

Boolean variable that is set based on the most recently received value of PMA:IS_SIGNAL indication (SIGNAL_OK). It is TRUE if the value was OK and FALSE if the value was FAIL.

slip_done

Boolean variable that is set to TRUE when the SLIP requested by the synchronization state diagram has been completed indicating that the next candidate FAS position can be tested.

test_fas

Boolean variable that is set to TRUE when a candidate FAS position is available for testing and FALSE when the FIND_1ST state is entered.

153.2.4.2 Functions**FAS_COMPARE**

This function compares the FAS with its expected value to determine if a valid frame alignment sequence has been detected and returns the result of the comparison using the variable `fas_match`. FAS_COMPARE is TRUE if `fas_valid` is TRUE, and `current_fecl` (interpreted with the most significant bit transmitted first) is equal to the `prev_fecl` plus 20 modulo 240. Otherwise, `fas_match` is FALSE.

SLIP

Causes the next candidate FAS position to be tested. The precise method for determining the next candidate FAS position is not specified and is implementation dependent. An implementation shall ensure that the FAS pattern can be detected in any possible position.

153.2.4.3 Counters

fas_bad_count

Counts the number of consecutive FASs that don't match the expected value for a given FEC lane.

fas_counter

Counts the 16 320 octets between the starting position of one FAS and the expected starting position of the next FAS on a FEC lane. This counter starts with a value of zero when start fas_counter is asserted (at the position of the first octet of a matched FAS pattern on a FEC lane), is incremented for each octet on that FEC lane, and reaches a terminal value of 16 320 at the expected position of the next FAS pattern on that FEC lane.

153.2.4.4 State diagrams

The SC-FEC shall implement 20 synchronization processes as shown in Figure 153-7. The synchronization process operates independently on each lane. The SC-FEC sublayer uses this process to detect the location of the frame alignment sequence in the received bit stream on each lane of the PMA service interface.

The SC-FEC shall implement the deskew process as shown in Figure 153-8.

IECNORM.COM : Click to view the full PDF of ISO/IEC/IEEE 8802-3:2021/Amd.13:2022

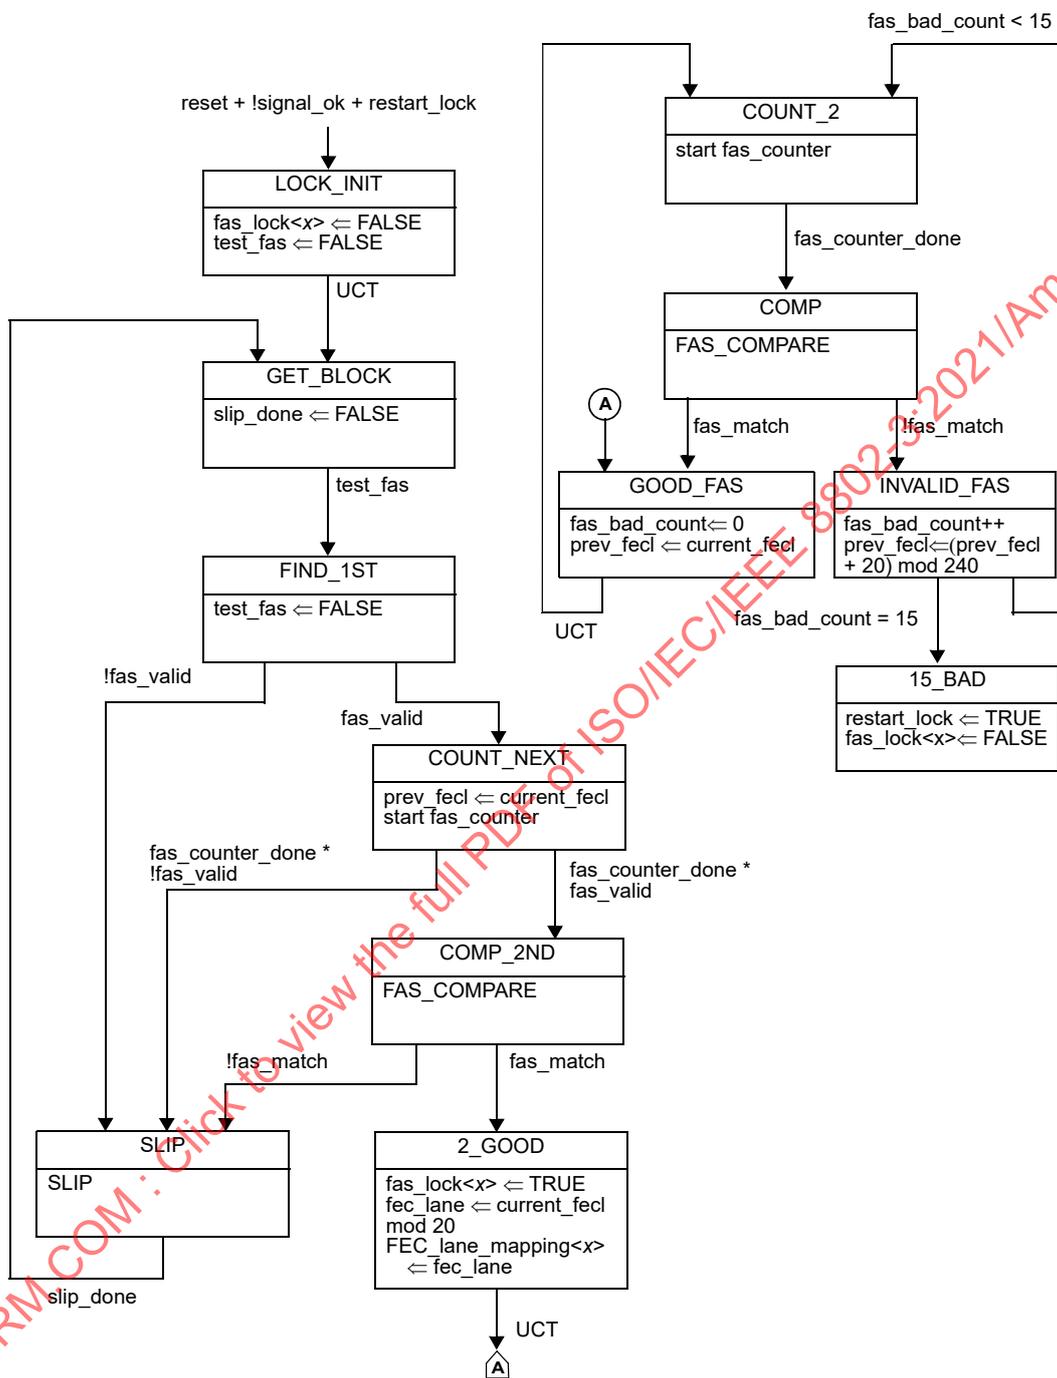


Figure 153-7—SC-FEC synchronization state diagram

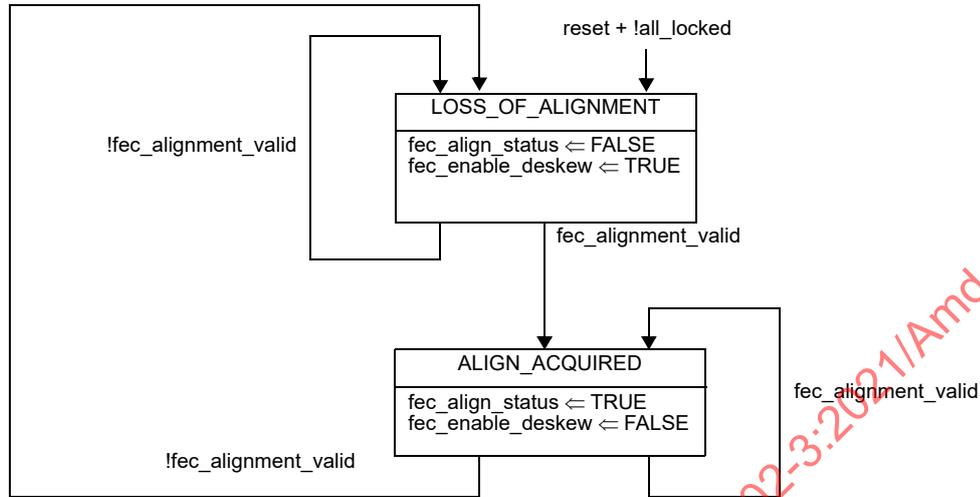


Figure 153–8—SC-FEC deskew state diagram

153.2.5 SC-FEC MDIO function mapping

The optional MDIO capability described in Clause 45 defines several registers that provide control and status information for and about the SC-FEC. If MDIO is implemented, it shall map MDIO status bits to SC-FEC status variables as shown in Table 153–2.

153.2.5.1 FEC_corrected_cw_counter

A corrected FEC codeword is a codeword that contained errors and was corrected.

The `FEC_corrected_cw_counter` is a 32-bit counter that counts once for each corrected FEC codeword processed when `fec_align_status` is TRUE. This variable is mapped to the registers defined in 45.2.1.186a (1.2276, 1.2277).

153.2.5.2 FEC_uncorrected_cw_counter

An uncorrected FEC codeword is a codeword that contains errors that were not corrected, including FEC codewords that may have been mis-corrected or not completely corrected.

The `FEC_uncorrected_cw_counter` is a 32-bit counter that counts once for each uncorrected FEC codeword processed when `fec_align_status` is TRUE. This variable is mapped to the registers defined in 45.2.1.186a (1.2278, 1.2279).

153.2.5.3 FEC_total_bits_counter

The `FEC_total_bits_counter` is a 64-bit counter that counts once for each bit processed by the SC-FEC decoder. This variable is mapped to the registers defined in 45.2.1.186a (1.2280, 1.2281, 1.2282, 1.2283). This may be used together with the `FEC_corrected_bits_counter` (see 153.2.5.4) as a measure of pre-FEC BER.

Table 153–2—MDIO/SC-FEC Status variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	FEC variable
SC-FEC FAS lock x , $x = 0$ to 19	SC-FEC alignment status 1 and 2 registers	1.2246.7:0 1.2247.11:0	fas_lock< x >
SC-FEC align status	SC-FEC alignment status 1 register	1.2246.12	fec_align_status
Lane x mapping, $x = 0$ to 19	SC-FEC lane mapping, lane x register	1.2250 through 1.2269	FEC_lane_mapping< x >
FEC corrected codewords	SC-FEC corrected codewords counter register	1.2276, 1.2277	FEC_corrected_cw_counter
FEC uncorrected codewords	SC-FEC uncorrected codewords counter register	1.2278, 1.2279	FEC_uncorrected_cw_counter
FEC total bits	SC-FEC total bits register	1.2280, 1.2281, 1.2282, 1.2283	FEC_total_bits_counter
FEC corrected bits	SC-FEC corrected bits register	1.2284, 1.2285, 1.2286, 1.2287	FEC_corrected_bits_counter

153.2.5.4 FEC_corrected_bits_counter

The FEC_corrected_bits_counter is a 64-bit counter that counts once for each bit corrected by the SC-FEC decoder. This variable is mapped to the registers defined in 45.2.1.186ao (1.2284, 1.2285, 1.2286, 1.2287). This may be used together with the FEC_total_bits_counter (see 153.2.5.3) to calculate pre-FEC BER.

153.3 100GBASE-ZR PMA sublayer**153.3.1 100GBASE-ZR PMA service interface**

This subclause specifies the services provided by the 100GBASE-ZR PMA sublayer. The 100GBASE-ZR PMA service interface is provided to allow the SC-FEC to transfer information to and from the 100GBASE-ZR PMD. The service interface is specified in an abstract manner and does not imply any particular implementation. The service interface primitives are summarized as follows:

PMA:IS_UNITDATA_ i .request
PMA:IS_UNITDATA_ i .indication
PMA:IS_SIGNAL.indication

The primitives are defined for $i = 0$ to 19. The SC-FEC continuously sends 20 parallel bit streams to the 100GBASE-ZR PMA sublayer, each at a nominal signaling rate of $(255/227) \times 4.97664$ Gb/s ± 20 ppm (~ 5.59049868 Gb/s). Likewise the 100GBASE-ZR PMA sublayer continuously sends 20 parallel bit streams to the SC-FEC sublayer, each at a nominal signaling rate of $(255/227) \times 4.97664$ Gb/s ± 20 ppm (~ 5.59049868 Gb/s).

153.3.2 Functions within the 100GBASE-ZR PMA sublayer

153.3.2.1 Functional block diagram

A functional diagram of the 100GBASE-ZR PMA sublayer is shown in Figure 153–9.

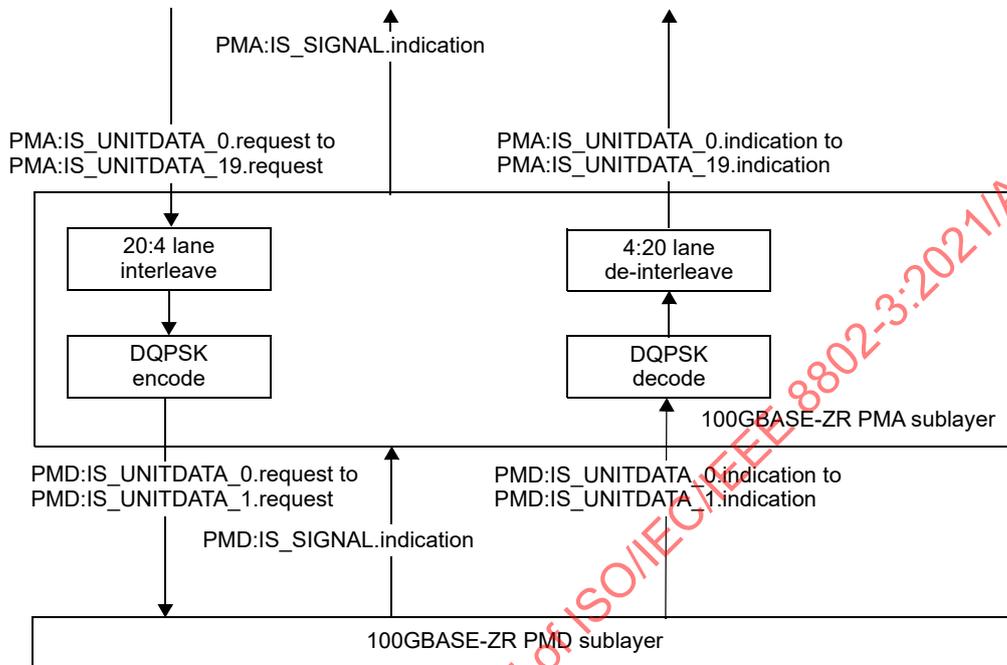


Figure 153–9—100GBASE-ZR PMA functional block diagram

153.3.2.2 Transmit function

153.3.2.2.1 Lane interleave

The lane interleave function bit-multiplexes groups of five lanes of the PMA service interface to form four logical bit streams operating at five times the service interface signaling rate. Each of the four bit streams constructed in this manner operates at a signaling rate of $(255/227) \times 24.8832 \text{ Gb/s} \pm 20 \text{ ppm}$ (~27.9525 Gb/s). The selection of which five lanes of the service interface is used to form each lane of the four-lane interface is arbitrary, but the combination of service interface lanes used to form each of the four lanes, and the order in which the bits are multiplexed from each of the input lanes, does not change while the link is in service. Each service interface lane is bit-multiplexed into exactly one lane of the four-lane interface.

NOTE—This four-lane interface format is referred to in ITU-T G.709 and ITU-T G.709.2 as OTL4.4.

153.3.2.2.2 DQPSK encode

The differential quadrature phase shift keying (DQPSK) encode function takes information from pairs of lanes (A, B) of the four-lane interface produced by the lane interleave function described in 153.3.2.2.1 and encodes them as a two streams of DQPSK symbols, one stream formed from each lane pair. The selection of the two lanes of the four-lane interface used to form each stream of DQPSK symbols is arbitrary, but does not change while the link is in service. The signaling rate of each stream of DQPSK symbols is $(255/227) \times 24.8832 \text{ GBd} \pm 20 \text{ ppm}$ (~27.9525 GBd).

With differential encoding, each DQPSK symbol can take one of four values 0 to 3. These four values are associated with the values of the pair of bits from the two lanes of the four-lane interface used to form each DQPSK symbol stream. The mapping of pairs of bits to DQPSK symbol values is given in Table 153–3.

Table 153–3—Bit mapping to DQPSK symbol values

Lane bit values A, B	DQPSK symbol value
00	0
10	1
11	2
01	3

153.3.2.3 Receive function

153.3.2.3.1 DQPSK decode

The 100GBASE-ZR PMA receives from the 100GBASE-ZR PMD two streams of DQPSK symbols at a signaling rate of $(255/227) \times 24.8832 \text{ GBd} \pm 20 \text{ ppm}$ ($\sim 27.9525 \text{ Gbd}$). The two streams of DQPSK symbols are converted to a four-lane digital interface as follows:

- 1) Each received DQPSK symbol is converted to a pair of bits as shown in Table 153–3.
- 2) One bit of each DQPSK symbol (A, B) is sent on a different lane of the four-lane digital interface. The assignment of each bit-position from the DQPSK symbols on the two DQPSK symbol streams to lanes of the four-lane digital interface is arbitrary, but does not change while the link is in service.

NOTE—This four-lane interface format is referred to in ITU-T G.709 and ITU-T G.709.2 as OTL4.4.

153.3.2.3.2 Lane de-interleave

Each lane of the four-lane digital signal produced by decoding of DQPSK symbols described in 153.3.2.3.1 is de-interleaved to five lanes of the 100GBASE-ZR PMA service interface by distributing successive bits round robin to each of the 100GBASE-ZR PMA service interface lanes. The assignment of bit positions on each of the lanes of the four-lane interface to the service interface lanes is arbitrary, but does not change while the link is in service.

153.4 Protocol implementation conformance statement (PICS) proforma for Clause 153, SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs⁴

153.4.1 Introduction

The supplier of a protocol implementation that is claimed to conform to Clause 153, SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs, shall complete the following protocol implementation conformance statement (PICS) proforma.

A detailed description of the symbols used in the PICS proforma, along with instructions for completing the PICS proforma, can be found in [Clause 21](#).

153.4.2 Identification

153.4.2.1 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).	

153.4.2.2 Protocol summary

Identification of protocol standard	IEEE Std 802.3ct-2021, Clause 153, SC-FEC and 100GBASE-ZR Physical Medium Attachment (PMA) sublayer for 100GBASE-ZR PHYs
Identification of amendments and corrigenda to this PICS proforma that have been completed as part of this PICS	
Have any Exception items been required? No [] Yes [] (See Clause 21 ; the answer Yes means that the implementation does not conform to IEEE Std 802.3ct-2021.)	

Date of Statement	
-------------------	--

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

153.4.3 Major capabilities/options

Item	Feature	Subclause	Value/Comment	Status	Support
DC	Delay constraints	153.2.2	Conforms to delay constraints specified in 153.2.2	M	Yes []
*MD	MDIO capability	153.2.5	Registers and interface supported	O	Yes [] No []

153.4.4 PICS proforma tables for SC-FEC sublayer for 100GBASE-ZR PHYs

153.4.4.1 Transmit function

Item	Feature	Subclause	Value/Comment	Status	Support
TF1	PCS lane Skew tolerance	153.2.3.2.2	Maximum Skew of 49 ps between PCS lanes and a maximum Skew Variation of 400 ps	M	Yes []
TF2	PCS lane reorder	153.2.3.2.3	Order the PCS lanes according to the PCS lane number	M	Yes []
TF3	GMP mapper	153.2.3.2.4	Create the SC-FEC frame and GMP map the deskewed and serialized PCS lanes into the payload area of the frame	M	Yes []
TF4	SC-FEC encoder	153.2.3.2.5	Compute and insert the SC-FEC parity into the SC-FEC frame	M	Yes []
TF5	Scrambler	153.2.3.2.6	Frame synchronous scrambler	M	Yes []
TF6	Lane distribution	153.2.3.2.7	Distribute SC-FEC frame to FEC lanes according to 153.2.3.2.7	M	Yes []

153.4.4.2 Receive function

Item	Feature	Subclause	Value/Comment	Status	Support
RF1	FEC lane Skew tolerance	153.2.3.3.1	Maximum Skew of 180 ns between FEC lanes and a maximum Skew Variation of 4 ns	M	Yes []
RF2	FEC lane reorder	153.2.3.3.2	Order the FEC lanes according to the FEC lane number	M	Yes []
RF3	De-scrambler	153.2.3.3.3	De-scramble according to 153.2.3.3.3	M	Yes []
RF4	SC-FEC decoder	153.2.3.3.4	Correct errors by decoding SC-FEC parity	M	Yes []
RF5	GMP demapper	153.2.3.3.5	Extract 66B encoded block stream from SC-FEC frame	M	Yes []
RF6	Block Alignment	153.2.3.3.6	Distribute blocks as PCS lanes to the FEC service interface	M	Yes []

153.4.4.3 State diagrams

Item	Feature	Subclause	Value/Comment	Status	Support
SD1	100GBASE-R block lock	153.2.3.2.1	Implements 20 block lock processes as depicted in Figure 82-12	M	Yes []
SD2	The SLIP function evaluates all possible bit positions	153.2.3.2.1		M	Yes []
SD3	100GBASE-R alignment marker lock	153.2.3.2.2	Implements 20 alignment marker lock processes as depicted in Figure 82-13	M	Yes []
SD4	The AM_SLIP function evaluates all possible blocks	153.2.3.2.2		M	Yes []
SD5	100GBASE-R PCS deskew state diagram	153.2.3.2.2	Meets the requirements of Figure 82-14	M	Yes []
SD6	SC-FEC synchronization state diagram	153.4.4.3	One instance per FEC lane per Figure 153-7	M	Yes []
SD7	FEC lane reorder	153.2.3.3.2	Order the FEC lanes according to the FEC lane number	M	Yes []

154. Physical Medium Dependent (PMD) sublayer and medium, type 100GBASE-ZR**154.1 Overview**

This clause specifies the 100GBASE-ZR PMD together with the associated medium, which is a single-mode fiber-based dense wavelength division multiplexing (DWDM) channel that may contain one or more optical amplifiers and is specified using a black link approach (see 154.6). The optical signal generated by this PMD type is modulated using a dual polarization differential quadrature phase shift keying (DP-DQPSK) format suitable for reception by a coherent optical receiver. When forming a complete Physical Layer, a PMD shall be connected to the 100GBASE-ZR PMA as shown in Table 154–1, to the medium through the MDI and optionally with the management functions that may be accessible through the management interface defined in Clause 45, or equivalent.

Table 154–1—Physical Layer clauses associated with the 100GBASE-ZR PMD

Associated clause	100GBASE-ZR
81—RS	Required
81—CGMII ^a	Optional
82—PCS	Required
83—100GBASE-R PMA	Optional
91—RS-FEC	Optional
83A—CAUI-10 C2C	Optional
83B—CAUI-10 C2M	Optional
83D—CAUI-4 C2C	Optional
83E—CAUI-4 C2M	Optional
152—Inverse RS-FEC	Conditional ^b
153—SC-FEC	Required
153—100GBASE-ZR PMA	Required
135—100GBASE-P PMA	Optional
135D—100GAUI-4 C2C	Optional
135E—100GAUI-4 C2M	Optional
135F—100GAUI-2 C2C	Optional
135G—100GAUI-2 C2M	Optional
78—Energy-Efficient Ethernet	Optional

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

^bClause 152 inverse RS-FEC mandatory when Clause 91 RS-FEC is present.

Figure 154–1 shows the relationship of the PMD and MDI (shown shaded) with other sublayers to the ISO/IEC Open System Interconnection (OSI) reference model. 100 Gigabit Ethernet is introduced in Clause 80 and the purpose of each PHY sublayer is summarized in 80.2.

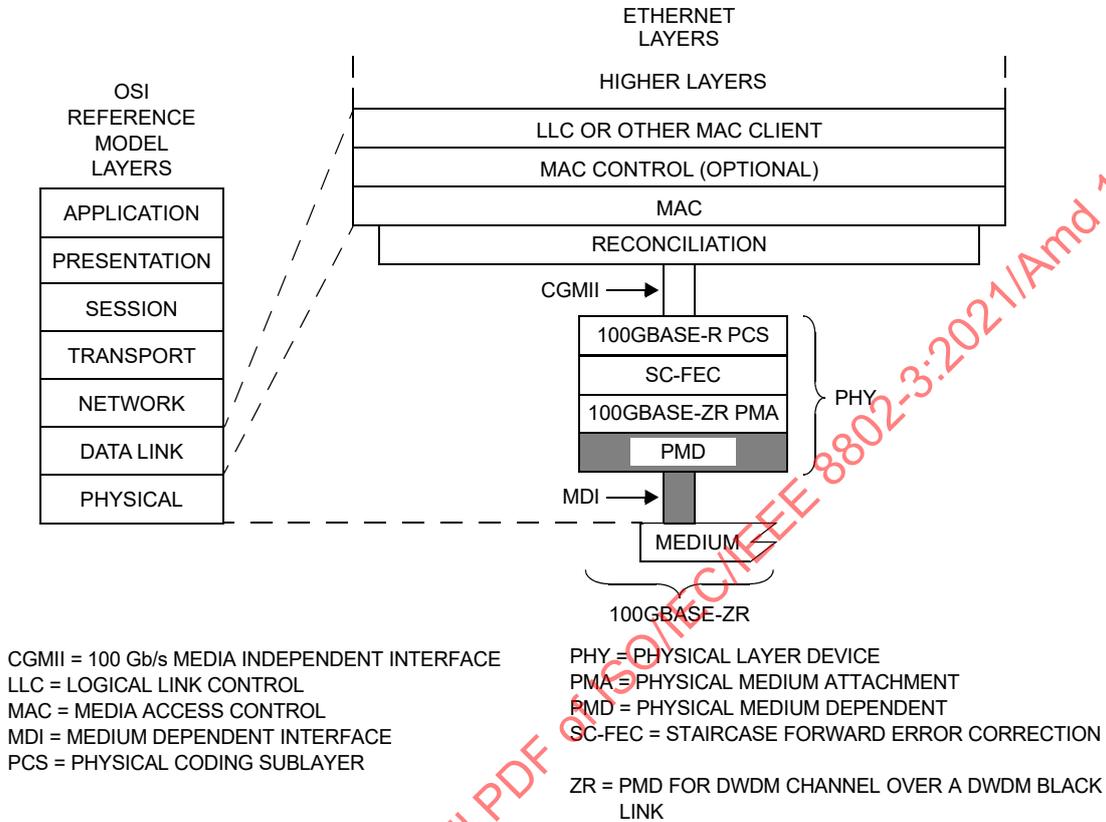


Figure 154–1—100GBASE-ZR PMD relationship to the ISO/IEC Open Systems Interconnection (OSI) reference model and IEEE 802.3 Ethernet model

100GBASE-ZR PHYs with the optional Energy-Efficient Ethernet (EEE) fast wake capability may enter the Low Power Idle (LPI) mode to conserve energy during periods of low link utilization (see Clause 78). The deep sleep mode of EEE is not supported.

154.1.1 Bit error ratio

The bit error ratio (BER) when processed by the PMA (Clause 153) shall be less than 4.62×10^{-3} provided that the error statistics are sufficiently random that this results in a frame loss ratio (see 1.4.275) of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap when additionally processed by the FEC (Clause 153) and PCS (Clause 82).

If the error statistics are not sufficiently random to meet this requirement, then the BER shall be less than that required to give a frame loss ratio of less than 6.2×10^{-10} for 64-octet frames with minimum interpacket gap.

154.2 Physical Medium Dependent (PMD) service interface

This subclause specifies the services provided by the 100GBASE-ZR PMD. The service interface for this PMD is described in an abstract manner and does not imply any particular implementation. The PMD service interface supports the exchange of encoded data between the PMA entity that resides just above the PMD, and the PMD entity. The PMD translates the encoded data to and from signals suitable for the specified medium.

The 100GBASE-ZR PMD service interface is an instance of the inter-sublayer service interface defined in 116.3, with two symbol streams, one for each polarization. The service interface primitives are summarized as follows:

```
PMD:IS_UNITDATA_i.request
PMD:IS_UNITDATA_i.indication
PMD:IS_SIGNAL.indication
```

The 100GBASE-ZR PMD has two parallel symbol streams, in which case $i = 0$ to 1.

In the transmit direction, the PMA continuously sends two parallel DQPSK symbol streams to the PMD, one stream per polarization, each with symbol values 0 to 3 and at a nominal signaling rate of $(255/227) \times 24.8832$ GBd (~ 27.9525 GBd), using the PMD:IS_UNITDATA_i.request primitive. The PMD then converts these streams of symbols into the appropriate signals on the MDI.

In the receive direction, the PMD continuously sends two parallel DQPSK symbol streams to the PMA, corresponding to the signals received from the MDI, each with symbol values 0 to 3, using the PMD:IS_UNITDATA_i.indication primitive, at a nominal signaling rate of $(255/227) \times 24.8832$ GBd (~ 27.9525 GBd).

The SIGNAL_OK parameter of the PMD:IS_SIGNAL.indication primitive corresponds to the variable SIGNAL_DETECT parameter as defined in 154.5.4. The SIGNAL_DETECT parameter takes a fixed value of OK.

NOTE—SIGNAL_DETECT = OK does not guarantee that the rx symbol parameters are known to be good. It is possible for a poor quality link to provide sufficient light for a SIGNAL_DETECT = OK indication and still not meet the BER defined in 154.1.1.

154.3 Delay and Skew

154.3.1 Delay constraints

The sum of the transmit and receive delays at one end of the link contributed by the 100GBASE-ZR PMD including 2 m of fiber in one direction shall be no more than 2048 bit times (4 pause_quanta or 20.48 ns). A description of overall system delay constraints and the definitions for bit times and pause_quanta can be found in 80.4 and its references.

154.3.2 Skew constraints

The Skew (relative delay) between the FEC lanes must be kept within limits so that the information on the FEC lanes can be reassembled by the FEC. The Skew Variation must also be limited to ensure that a given FEC lane always traverses the same physical lane. Skew and Skew Variation are defined in 80.5 and specified at the points SP0 to SP7 shown in Figure 80–8.

If the PMD service interface is physically instantiated so that the Skew at SP2 can be measured, then the Skew at SP2 is limited to 43 ns and the Skew Variation at SP2 is limited to 400 ps.

The Skew at SP3 (the transmitter MDI) shall be less than 54 ns and the Skew Variation at SP3 shall be less than 600 ps.

The Skew at SP4 (the receiver MDI) shall be less than 134 ns and the Skew Variation at SP4 shall be less than 3.4 ns.

If the PMD service interface is physically instantiated so that the Skew at SP5 can be measured, then the Skew at SP5 shall be less than 145 ns and the Skew Variation at SP5 shall be less than 3.6 ns.

For more information on Skew and Skew Variation, see 80.5.

154.4 PMD MDIO function mapping

The optional MDIO capability described in Clause 45 defines several variables that may provide control and status information for and about the PMD. If the MDIO interface is implemented, the mapping of MDIO control variables to PMD control variables shall be as shown in Table 154–2 and the mapping of MDIO status variables to PMD status variables shall be as shown in Table 154–3.

Table 154–2—MDIO/PMD control variable mapping

MDIO control variable	PMA/PMD register name	Register/bit number	PMD control variable
Reset	PMA/PMD control 1 register	1.0.15	PMD_reset
Global PMD transmit disable	PMD transmit disable register	1.9.0	PMD_global_transmit_disable
Tx optical channel index	Tx optical channel control register	1.800.5:0	Tx_optical_channel_index
Rx optical channel index	Rx optical channel control register	1.820.5:0	Rx_optical_channel_index

Table 154–3—MDIO/PMD status variable mapping

MDIO status variable	PMA/PMD register name	Register/bit number	PMD status variable
Fault	PMA/PMD status 1 register	1.1.7	PMD_fault
Transmit fault	PMA/PMD status 2 register	1.8.11	PMD_transmit_fault
Receive fault	PMA/PMD status 2 register	1.8.10	PMD_receive_fault
Global PMD receive signal detect	PMD receive signal detect register	1.10.0	PMD_global_signal_detect
Tx index ability 0 to Tx index ability 47	Tx optical channel ability 1 register to Tx optical channel ability 3 register	1.801 to 1.803	Tx_index_ability_0 to Tx_index_ability_47
Tx Rx different optical channel ability	Rx optical channel control register	1.820.15	Tx_Rx_diff_opt_chan_ability
Rx index ability 0 to Rx index ability 47	Rx optical channel ability 1 register to Rx optical channel ability 3 register	1.821 to 1.823	Rx_index_ability_0 to Rx_index_ability_47