

INTERNATIONAL IEEE Std 1588™ STANDARD

Precision Clock Synchronization Protocol for Networked Measurement and Control Systems



THIS PUBLICATION IS COPYRIGHT PROTECTED
Copyright © 2019 IEEE

All rights reserved. IEEE is a registered trademark in the U.S. Patent & Trademark Office, owned by the Institute of Electrical and Electronics Engineers, Inc. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from the IEC Central Office. Any questions about IEEE copyright should be addressed to the IEEE. Enquiries about obtaining additional rights to this publication and other information requests should be addressed to the IEC or your local IEC member National Committee.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland
Tel.: +41 22 919 02 11
info@iec.ch
www.iec.ch

Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue
New York, NY 10016-5997
United States of America
stds.info@ieee.org
www.ieee.org

About the IEC

The International Electrotechnical Commission (IEC) is the leading global organization that prepares and publishes International Standards for all electrical, electronic and related technologies.

About IEC publications

The technical content of IEC publications is kept under constant review by the IEC. Please make sure that you have the latest edition, a corrigendum or an amendment might have been published.

IEC publications search - webstore.iec.ch/advsearchform

The advanced search enables to find IEC publications by a variety of criteria (reference number, text, technical committee, ...). It also gives information on projects, replaced and withdrawn publications.

IEC Just Published - webstore.iec.ch/justpublished

Stay up to date on all new IEC publications. Just Published details all new publications released. Available online and once a month by email.

IEC Customer Service Centre - webstore.iec.ch/csc

If you wish to give us your feedback on this publication or need further assistance, please contact the Customer Service Centre: sales@iec.ch.

IEC online collection - oc.iec.ch

Discover our powerful search engine and read freely all the publications previews. With a subscription you will always have access to up to date content tailored to your needs.

Electropedia - www.electropedia.org

The world's leading online dictionary on electrotechnology, containing more than 22 000 terminological entries in English and French, with equivalent terms in 18 additional languages. Also known as the International Electrotechnical Vocabulary (IEV) online.

INTERNATIONAL IEEE Std 1588™ STANDARD

Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 25.040.40; 35.110; 35.240.50

IEEE (PDF) STD 24776

ISBN 978-2-8322-9811-4
ISBN 978-1-5044-7667-6

Warning! Make sure that you obtained this publication from an authorized distributor.

Contents

1. Scope.....	14
2. Normative references	15
3. Definitions, acronyms, and abbreviations.....	16
3.1 Definitions	16
3.2 Acronyms and abbreviations.....	24
4. Conventions.....	26
4.1 Descriptive lexical form syntax.....	26
4.2 Word usage	27
4.3 Behavioral specification notation	28
5. Data types and on-the-wire formats	30
5.1 General	30
5.2 Primitive data type specifications.....	30
5.3 Derived data type specifications.....	31
5.4 On-the-wire formats	35
6. Clock synchronization model	36
6.1 General requirements on implementations.....	36
6.2 Principal assumptions about the network and implementation recommendations.....	39
6.3 PTP Networks	39
6.4 PTP message classes.....	40
6.5 PTP device types	41
6.6 Synchronization overview.....	53
6.7 PTP communications overview.....	65
7. Characterization of PTP entities	69
7.1 Domains.....	69
7.2 Timescales used in PTP	72
7.3 PTP communications.....	74
7.4 PTP communication media	81
7.5 PTP Ports	84
7.6 PTP Instance characterization	89
7.7 PTP timing characterization.....	103
8. PTP data sets.....	105
8.1 General specifications for data set members.....	105
8.2 Data sets for PTP Instances.....	113
8.3 Data sets for Transparent Clocks.....	141
8.4 commonMeanLinkDelayService data sets	144
9. PTP for Ordinary Clocks and Boundary Clocks	144
9.1 General protocol requirements for PTP Ordinary Clocks and Boundary Clocks	144
9.2 State protocol	145
9.3 Best master clock algorithms	154
9.4 Grandmaster PTP Instance timePropertiesDS updates	164
9.5 PTP message processing semantics	166
9.6 Changes in the PTP Instance.....	182
10. PTP for Transparent Clocks.....	182
10.1 Requirements for both end-to-end and peer-to-peer Transparent Clocks.....	182
10.2 End-to-end Transparent Clock requirements.....	182
10.3 Peer-to-peer Transparent Clock requirements	191
11. Clock offset, path delay, residence time, and asymmetry corrections	195

11.1 General specifications.....	195
11.2 Computation of <offsetFromMaster> in Ordinary Clocks and Boundary Clocks	196
11.3 Delay request-response mechanism for Ordinary Clocks and Boundary Clocks.....	197
11.4 Peer-to-peer delay mechanism	199
11.5 MDMI interface and Special Ports	204
12. Synchronization and syntonization of clocks.....	214
12.1 Clock adjustments	214
12.2 Syntonization	214
12.3 Synchronization.....	215
13. PTP message formats	216
13.1 General.....	216
13.2 General PTP message format requirements.....	216
13.3 Header.....	216
13.4 Suffix.....	222
13.5 Announce message	222
13.6 Sync and Delay_Req messages	224
13.7 Follow_Up message	224
13.8 Delay_Resp message	225
13.9 Pdelay_Req message	225
13.10 Pdelay_Resp message	226
13.11 Pdelay_Resp_Follow_Up message.....	226
13.12 Signaling message	227
13.13 PTP management message.....	228
14. TLV entity specifications	228
14.1 General requirements.....	228
14.2 Propagation of TLVs through Boundary Clocks	230
14.3 Vendor and standard organization extension TLVs.....	231
14.4 PAD TLV (optional).....	233
15. PTP management messages (optional)	234
15.1 General.....	234
15.2 PTP management mechanism	234
15.3 Processing of PTP management messages.....	234
15.4 PTP management message format.....	235
15.5 Management TLVs.....	238
16. General optional features.....	272
16.1 Unicast message negotiation (optional)	272
16.2 Path trace (optional).....	280
16.3 Alternate timescale offsets (optional)	282
16.4 Holdover upgrade (optional).....	289
16.5 Isolation of PTP Instances running under profiles specified by different standards organizations (optional)	290
16.6 Common Mean Link Delay Service (optional).....	291
16.7 Configurable correction of timestamps (optional).....	298
16.8 Calculation of the <delayAsymmetry> for certain media (optional)	299
16.9 Mixed multicast/unicast operation (optional).....	301
16.10 Cumulative frequency transfer method for synchronizing clocks (optional).....	304
16.11 Slave Event Monitoring (optional)	308
16.12 Enhanced synchronization accuracy metrics (optional)	317
16.13 Message Length Extension (optional).....	323
16.14 PTP integrated security mechanism (optional).....	324
17. State configuration options.....	338
17.1 General.....	338
17.2 Grandmaster clusters (optional).....	338

17.3 Alternate master (optional).....	340
17.4 Unicast discovery (optional)	343
17.5 Acceptable master table (optional)	345
17.6 Mechanism for external configuration of a PTP Instance's PTP Port state (optional).....	347
17.7 Reduced state sets and use of the <foreignMasterList> feature (optional).....	353
18. Interactions between PTP Instances in different PTP domains	354
18.1 General specifications.....	354
18.2 Interfaces enabling interdomain interactions.....	355
19. Compatibility of this edition with earlier and future editions	355
19.1 General.....	355
19.2 Compatibility between version 2 and future versions.....	356
19.3 Compatibility with IEEE Std 1588-2002	356
19.4 Compatibility between the PTP Instance conformant to this edition and the implementations conformant to IEEE Std 1588-2008	356
20. Conformance	361
20.1 Conformance objective	361
20.2 PTP conformance requirements.....	361
20.3 PTP Profiles	362
Annex A (informative) Using the Precision Time Protocol (PTP).....	365
A.1 Overview.....	365
A.2 Physical layout	366
A.3 Logical layout.....	366
A.4 Component issues.....	367
A.5 Local implementation issues	368
A.6 System implementation issues.....	370
A.7 Guidelines to achieve optimal performance	371
A.8 Recommendations to aid in conformance testing.....	371
A.9 Recommendation for implementations in unicast networks or networks with non-PTP bridges and routers	372
Annex B (informative) Timescales and epochs in PTP.....	375
B.1 General considerations.....	375
B.2 UTC, TAI and the PTP epoch and timescale updates	375
B.3 Standard time sources	378
B.4 Meaning and uses of the attributes of the timePropertiesDS data set.....	379
Annex C (normative) Transport of PTP over User Datagram Protocol over Internet Protocol Version 4.....	382
C.1 General.....	382
C.2 UDP port numbers	382
C.3 IPv4 multicast addresses	382
C.4 sdoId field values.....	383
C.5 Optional values.....	384
C.6 IPv4 Options.....	384
C.7 Protocol addresses	384
Annex D (normative) Transport of PTP over User Datagram Protocol over Internet Protocol Version 6.....	385
D.1 General.....	385
D.2 UDP port numbers.....	385
D.3 IPv6 multicast addresses	386
D.4 Optional values.....	386
D.5 Protocol addresses	386
Annex E (normative) Transport of PTP over IEEE 802.3 transports	387
E.1 General.....	387
E.2 Ethertype	387
E.3 Multicast media access control (MAC) addresses.....	387

E.4 majorSdoId field values.....	388
E.5 Optional values.....	388
E.6 Protocol addresses.....	388
Annex F (normative) Transport of PTP over DeviceNET.....	389
F.1 Protocol.....	389
F.2 message timestamp point.....	389
F.3 clockIdentity.....	389
F.4 PTP message formats.....	389
F.5 DeviceNet addressing for PTP.....	390
Annex G (normative) Transport of PTP over ControlNET.....	391
G.1 Protocol.....	391
G.2 clockIdentity.....	391
G.3 PTP message formats.....	391
G.4 ControlNet addressing for PTP.....	391
Annex H (normative) Transport of PTP over IEC 61158 Type 10.....	392
H.1 Background.....	392
H.2 Message specification.....	393
H.3 DLPDU of the IEC 61158 TYPE10.....	394
H.4 Encoding specifications.....	395
Annex I (normative) Default PTP Profiles.....	398
I.1 General.....	398
I.2 General requirements.....	398
I.3 Delay Request-Response Default PTP Profile.....	398
I.4 Peer-to-Peer Default PTP Profile.....	400
I.5 High-Accuracy Delay Request-Response Default PTP Profile.....	401
Annex J (normative) Performance monitoring options (optional).....	407
J.1 General.....	407
J.2 Timestamp monitoring.....	407
J.3 Additional parameters.....	410
J.4 Record data types.....	412
J.5 Data sets for performance monitoring.....	414
Annex K (informative) Suppression of rogue Announce messages.....	419
K.1 Example—Star topology.....	420
K.2 Example—PTP Network with a single loop with an odd number of PTP Instances in the loop	422
K.3 Example—More complex single loop PTP Network.....	423
K.4 Example—Linear chain.....	423
Annex L (normative) Layer-1 based synchronization performance enhancement (optional).....	425
L.1 General.....	425
L.2 Basic terms.....	426
L.3 Link Reference Model.....	427
L.4 L1Sync port characteristics.....	429
L.5 L1Sync data sets.....	431
L.6 L1Sync message exchange.....	434
L.7 L1Sync port operation specification.....	437
L.8 Optional parameters (option within this option).....	439
L.9 Link verification using Signaling messages (informative).....	444
Annex M (informative) Sub-nanosecond synchronization using the High Accuracy Default PTP Profile.....	445
M.1 General.....	445
M.2 Frequency loopback.....	445
M.3 Timestamping precision.....	447
M.4 Timestamping accuracy.....	448

M.5 Medium and its asymmetry	449
M.6 Timing characteristics	450
Annex N (informative) Calibration procedures	452
N.1 General.....	452
N.2 Theoretical background.....	455
N.3 Assumptions and requirements.....	457
N.4 Calibration procedures.....	459
Annex O (informative) Example inter-domain interactions	466
O.1 General.....	466
O.2 Sourcing timing to multiple domains.....	466
O.3 Providing timing to users (sinks) from multiple independent domains.....	467
O.4 Transferring time from PTP domain A to PTP domain B	468
O.5 Example use for external configuration of port state	473
Annex P (informative) Security.....	474
P.1 Overview, assumptions, and approach	474
P.2 Multipronged approach—detailed definition	474
Annex Q (informative) Bibliography.....	495
Annex R (informative) IEEE List of participants	498

PRECISION CLOCK SYNCHRONIZATION PROTOCOL FOR NETWORKED MEASUREMENT AND CONTROL SYSTEMS

FOREWORD

- 1) The International Electrotechnical Commission (IEC) is a worldwide organization for standardization comprising all national electrotechnical committees (IEC National Committees). The object of IEC is to promote international co-operation on all questions concerning standardization in the electrical and electronic fields. To this end and in addition to other activities, IEC publishes International Standards, Technical Specifications, Technical Reports, Publicly Available Specifications (PAS) and Guides (hereafter referred to as "IEC document(s)"). Their preparation is entrusted to technical committees; any IEC National Committee interested in the subject dealt with may participate in this preparatory work. International, governmental and non-governmental organizations liaising with the IEC also participate in this preparation.

IEEE Standards documents are developed within IEEE Societies and Standards Coordinating Committees of the IEEE Standards Association (IEEE SA) Standards Board. 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 IEEE and serve without compensation. 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 contained in its standards. Use of IEEE Standards documents is wholly voluntary. *IEEE documents are made available for use subject to important notices and legal disclaimers (see <http://standards.ieee.org/ipr/disclaimers.html> for more information).*

IEC collaborates closely with IEEE in accordance with conditions determined by agreement between the two organizations. This Dual Logo International Standard was jointly developed by the IEC and IEEE under the terms of that agreement.

- 2) The formal decisions of IEC on technical matters express, as nearly as possible, an international consensus of opinion on the relevant subjects since each technical committee has representation from all interested IEC National Committees. The formal decisions of IEEE on technical matters, once consensus within IEEE Societies and Standards Coordinating Committees has been reached, is determined by a balanced ballot of materially interested parties who indicate interest in reviewing the proposed standard. Final approval of the IEEE standards document is given by the IEEE Standards Association (IEEE SA) Standards Board.
- 3) IEC/IEEE Publications have the form of recommendations for international use and are accepted by IEC National Committees/IEEE Societies in that sense. While all reasonable efforts are made to ensure that the technical content of IEC/IEEE Publications is accurate, IEC or IEEE cannot be held responsible for the way in which they are used or for any misinterpretation by any end user.
- 4) In order to promote international uniformity, IEC National Committees undertake to apply IEC Publications (including IEC/IEEE Publications) transparently to the maximum extent possible in their national and regional publications. Any divergence between any IEC/IEEE Publication and the corresponding national or regional publication shall be clearly indicated in the latter.
- 5) IEC and IEEE do not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. IEC and IEEE are not responsible for any services carried out by independent certification bodies.
- 6) All users should ensure that they have the latest edition of this publication.
- 7) No liability shall attach to IEC or IEEE or their directors, employees, servants or agents including individual experts and members of technical committees and IEC National Committees, or volunteers of IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE SA) Standards Board, for any personal injury, property damage or other damage of any nature whatsoever, whether direct or indirect, or for costs (including legal fees) and expenses arising out of the publication, use of, or reliance upon, this IEC/IEEE Publication or any other IEC or IEEE Publications.
- 8) Attention is drawn to the normative references cited in this publication. Use of the referenced publications is indispensable for the correct application of this publication.
- 9) Attention is drawn to the possibility that implementation of this IEC/IEEE Publication may require use of material covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. IEC or IEEE shall not be held responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patent Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this standard are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility.

IEC 61588/IEEE Std 1588 was processed through IEC technical committee 65C: Industrial networks, under the IEC/IEEE Dual Logo Agreement. It is an International Standard.

The text of this International Standard is based on the following documents:

IEEE Std	FDIS	Report on voting
1588 (2019)	65C/1084/FDIS	65C/1091/RVD

Full information on the voting for its approval can be found in the report on voting indicated in the above table.

The language used for the development of this International Standard is English.

The IEC Technical Committee and IEEE Technical Committee have decided that the contents of this document will remain unchanged until the stability date indicated on the IEC website under webstore.iec.ch in the data related to the specific document. At this date, the document will be

- reconfirmed,
- withdrawn,
- replaced by a revised edition, or
- amended.

Abstract: In this standard, a protocol is defined that provides precise synchronization of clocks in packet-based networked systems. Synchronization of clocks can be achieved in heterogeneous systems that include clocks of different inherent precision, resolution, and stability. The protocol supports synchronization accuracy and precision in the sub-microsecond range with minimal network and local computing resources. Customization is supported by means of profiles. The protocol includes default profiles that permit simple systems to be installed and operated without the need for user management. Sub-nanosecond time transfer accuracy can be achieved in a properly designed network.

Keywords: Boundary Clock, clock, Grandmaster Clock, IEEE 1588™, management, Ordinary Clock, security, synchronization, Transparent Clock

Important Notices and Disclaimers Concerning IEEE Standards Documents

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

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

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

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

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

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

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

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

Translations

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

Official statements

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

Comments on standards

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

Comments on standards should be submitted to the following address:

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

Laws and regulations

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

Copyrights

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

Updating of IEEE Standards documents

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

Every IEEE standard is subjected to review at least every 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 at <https://ieeexplore.ieee.org> or contact IEEE at the address listed previously. For more information about the IEEE SA or IEEE's standards development process, visit the IEEE SA Website at <http://standards.ieee.org>.

Errata

Errata, if any, for IEEE standards can be accessed via <https://standards.ieee.org/standard/index.html>. 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: <https://ieeexplore.ieee.org/browse/standards/collection/ieee/>. Users are encouraged to periodically check for errata.

Patents

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

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

Introduction

This introduction is not part of IEEE Std 1588-2019, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.

This standard defines a protocol that provides precise synchronization of clocks in packet-based networked systems. The Precision Time Protocol (PTP) generates a master–slave relationship among the PTP Instances in the system. The clocks in all PTP Instances ultimately derive their time from a clock known as the “Grandmaster Clock.” In its basic form, this protocol is intended to be administration free.

IEEE Std 1588-2019 includes content that was not present in IEEE Std 1588-2008. Similarly some content that was present in the IEEE Std 1588-2008 is not in IEEE Std 1588-2019. The following Annexes in 1588-2008 are not present in IEEE Std 1588-2019:

- Annex C (informative) Examples of residence and asymmetry corrections
- Annex K (informative) Security protocol (experimental)
- Annex L (informative) Transport of cumulative frequency scale factor offset (experimental)

Acknowledgments

The working group would like to acknowledge Dieter Sibold and Steffen Fries for their contributions to the integrated security mechanism.

IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems

1. Scope

This standard defines a network protocol, the Precision Time Protocol (PTP), enabling accurate and precise synchronization of the real-time clocks of devices in networked distributed systems. The protocol is applicable to systems where devices communicate via networks, including Ethernet. The standard allows multicast communication, unicast communication or both. The standard specifies requirements for mapping the protocol to specific network implementations and defines such mappings, including User Datagram Protocol (UDP)/Internet Protocol (IP versions 4 and 6), and layer-2 IEEE 802.3 Ethernet.

The protocol enables heterogeneous systems that include clocks of various inherent precision, resolution, and stability to synchronize to a grandmaster clock. The protocol supports synchronization in the sub-microsecond range with minimal network bandwidth and local clock computing resources. The protocol enhances support for synchronization to better than 1 nanosecond. The protocol specifies how corrections for path asymmetry are made, if the asymmetry values are known. The grandmaster can be synchronized to a source of time external to the system, if time traceable to international standards or other source of time is required. The protocol provides information for devices to compute Coordinated Universal Time (UTC) from the protocol distributed time, if the grandmaster is traceable to international standards and is able to access pending leap-second changes. Options are also provided to allow end devices to compute other time scales from the protocol distributed time scale.

The protocol defines timing domains in which system timing is consistent. The protocol establishes the timing topology. The default behavior of the protocol allows simple systems to be installed and operated without requiring the administrative attention of users to determine the system timing topology.

The standard defines all needed data types, message formats, required computations, internal states, the behavior of devices with respect to transmitting, receiving, and processing protocol communications. The standard provides for the management of protocol artifacts in devices. The standard defines formal mechanisms for message extensions and the requirements for profiles that allow customization for specific application domains.

The standard defines conformance requirements. Optional specifications are provided for protocol security. This standard documents conditions under which this standard is backward compatible with IEEE 1588-2008.

2. Normative references

The following referenced documents are indispensable for the application of this document (i.e., they must be understood and used, so each referenced document is cited in text and its relationship to this document is explained). For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments or corrigenda) applies.

FIPS PUB 180-4:2015, Secure Hash Standard.¹

FIPS PUB 198-1:2008, The Keyed-Hash Message Authentication Code (HMAC).

IEC 61158-5-10:2007, Industrial communication networks—Fieldbus specifications—Part 5-10: Application layer service definition—Type 10 elements.²

IEC 61158-6-10:2007, Industrial communication networks—Fieldbus specifications—Part 6-10: Application layer protocol specification—Type 10 elements.

IEC 61784-1:2007, Industrial communication networks—Profiles—Part 1: Fieldbus profiles.

IEC 61784-2:2007, Industrial communications networks—Profiles—Part 2: Additional fieldbus profiles for real-time networks based on ISO/IEC 8802-3.

IEC 62026-3:2008, Low-voltage switchgear and controlgear—Controller-device interfaces (CDIs)—Part 3: DeviceNet.

IEEE Std 754™, IEEE Standard for Binary Floating-Point Arithmetic.^{3, 4}

IEEE Std 802®, IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture.

IEEE Std 802.1AS™-2011, IEEE Standard for Local and Metropolitan Area Networks—Timing and Synchronization for Time-Sensitive Applications in Bridged Local Area Networks.

IEEE Std. 802c™, Standard for Local and Metropolitan Area Networks—Overview and Architecture Amendment: Local Medium Access Control (MAC) Address Usage.

IEEE Std 802.1Q™ -2014, IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks.

IEEE Std 802.3™, IEEE Standard for Ethernet.

IEEE Std. 802.11™, IEEE Standard for Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.

IERS Bulletin C.⁵

IETF RFC 768 (1980), User Datagram Protocol, Postel, J., ed.⁶

IETF RFC 791 (1981), Internet Protocol, Postel, J., ed.

¹ FIPS publications are available from the National Technical Information Service, U. S. Department of Commerce (<http://www.ntis.org/>).

² IEC publications are available from the International Electrotechnical Commission (<http://www.iec.ch>) and the American National Standards Institute (<http://www.ansi.org/>).

³ IEEE publications are available from The Institute of Electrical and Electronics Engineers (<http://standards.ieee.org/>).

⁴ The IEEE standards or products referred to in Clause 2 are trademarks owned by The Institute of Electrical and Electronics Engineers, Incorporated.

⁵ IERS publications are available from the International Earth Rotation and Reference Systems Services (<https://www.iers.org/>).

⁶ IETF publications are available from the Internet Engineering Task Force (<https://www.ietf.org/>).

IETF RFC 1624 (1994), Computation of the Internet Checksum via Incremental Update, Rijssinghani, A., ed.

IETF RFC 4868 (2007), Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec, Kelly, S., and Frankel, S., eds.

IETF RFC 8200 (2017), Internet Protocol, Version 6 (IPv6) Specification, Deering, S., and Hinden, R., eds.

ISO/IEC 10646:2003, Information technology—Universal Multiple-Octet Coded Character Set (UCS)—Part 1: Architecture and Basic Multilingual Plane.⁷

3. Definitions, acronyms, and abbreviations

3.1 Definitions

For the purposes of this document, the following terms and definitions apply.

3.1.1 accuracy: The mean over an ensemble of measurements of the time or frequency error between the clock under test and a reference clock.

NOTE 1—Line “B” in Figure 1 represents the error in the measured mean value with respect to the reference, that is, the accuracy. The width of the curve of ensemble measurements, “C” in Figure 1, represents the measurement precision.⁸




Figure 1—Histogram of an ensemble of measurements

NOTE 2—Different industries define "accuracy" in differing ways; however, for the discussion of an IEEE 1588 Grandmaster Clock, a smaller mean or maximum error from the reference time scale will be considered more accurate (whether the mean error or the maximum error is more relevant depends on the application).

3.1.2 atomic process: A process in which all the values of all inputs to the process are not permitted to change until all the results of the process are instantiated, and the outputs of the process are not visible to other processes until the processing of each output is complete.

⁷ISO publications are available from the International Organization for Standardization (<http://www.iso.org/>) and the American National Standards Institute (<http://www.ansi.org/>).

⁸Notes in text, tables, and figures of a standard are given for information only and do not contain requirements needed to implement this standard.

3.1.3 Boundary Clock: A PTP Instance that has multiple PTP Ports in a domain and maintains the timescale used in the domain. Within a domain, it may serve as the source of time to other PTP Instances, that is, be a Master Clock, and it can in addition synchronize to another Boundary Clock or Ordinary Clock, that is, be a Slave Clock.

3.1.4 clock: A device that can provide a measurement of the passage of time since a defined epoch.

NOTE 1—There are only two types of clocks specified in this standard: the Local PTP Clock and the Local Clock.

NOTE 2—For Local PTP Clocks associated with Ordinary Clocks and Boundary Clocks that are properly synchronized, the epoch is the epoch of the timescale in use. In the case of a Local Clock associated with a Transparent Clock, the epoch is locally defined and not necessarily aligned with the timescale in use.

NOTE 3—A clock provides time at desired moments of the timescale it maintains. Time is obtained either:

- *Physically:* in this type of clock, the time is modeled using a clock signal and a time counter that is driven by the clock signal;
- *Mathematically:* in this type of clock, the time is generated by a model that describes the relation of this clock to another clock (e.g., to a physical clock in a different timescale). The model enables the calculation of the time of the clock from the time of the other clock.

NOTE 4—The word “clock” appears many times in this standard modified as Boundary Clock, Ordinary Clock, Transparent Clock, or Grandmaster Clock. When used this way the combination, e.g. Boundary Clock, is not a clock as defined here but rather contains a clock. For historical reasons, these terms have been retained.

3.1.5 clock signal: A physical signal that has periodic events. The periodic events mark the significant instants at which a time counter is incremented. The clock signal is characterized by its frequency and phase.

3.1.6 default: When applied to attribute values and options of a PTP Instance, “default” means the configuration of a PTP Instance as it is delivered from the manufacturer.

3.1.7 device: An entity implementing some functionality, for example, a clock, a Boundary Clock, an encoder, and a port. *See also:* **PTP Instance**.

3.1.8 Direct PTP Link: A direct physical link with no intervening network elements between two PTP Ports of Boundary Clocks, Ordinary Clocks, and/or peer-to-peer Transparent Clocks.

NOTE 1—A direct physical link, where at least one end point is a PTP Port of an end-to-end Transparent Clock, is not a Direct PTP Link.

NOTE 2—A Direct PTP Link is a PTP Communication Path if the delay request-response mechanism is used. It is a PTP Link if the peer-to-peer delay mechanism is used, and is at the same time a PTP Communication Path if it is between two Boundary Clocks, two Ordinary Clocks, or a Boundary Clock and an Ordinary Clock.

3.1.9 domain (or PTP Domain): A logical grouping of PTP Instances using PTP to ensure that all Local PTP Clocks in the grouping are synchronized to the Grandmaster Clock of the domain, but are not necessarily synchronized to the Local PTP Clocks in another domain.

3.1.10 end-to-end port: A PTP Port that is configured to use the delay request-response mechanism.

3.1.11 end-to-end Transparent Clock: A Transparent Clock that supports the use of the delay request-response mechanism between a PTP Port in the MASTER state and a PTP Port in the SLAVE state in the same domain.

3.1.12 epoch: The origin of a timescale.

3.1.13 event: An abstraction of the mechanism by which signals or conditions are generated and represented. In this abstraction, the aspects of interest of the signals are conditions that occur at discrete instants of time.

3.1.14 foreign master: If an Announce message is received on PTP Port a of PTP Instance A from PTP Port b of PTP Instance B in the same PTP Network, and PTP Port b is not the current Master of PTP Port a, then PTP Port b is considered a foreign master by PTP Port a.

3.1.15 fractional frequency offset: The fractional frequency offset (FFO) between a measured frequency and a reference frequency is defined as follows:

$$\text{FFO} = \frac{(\text{FO} - \text{FR})}{\text{FR}} \quad (1)$$

where

FO is the measured frequency
FR is the reference frequency

3.1.16 frequency offset: The frequency offset between a measured frequency and a reference frequency is defined as follows:

$$\text{Frequency Offset} = \text{FO} - \text{FR} \quad (2)$$

where

FO is the measured frequency,
FR is the reference frequency.

3.1.17 Grandmaster Clock: In the context of a single PTP domain, the Local PTP Clock of an Ordinary Clock or a Boundary Clock that is the source of time to which all other Local PTP Clocks in the domain are synchronized.

3.1.18 Grandmaster PTP Instance: A PTP Instance containing the Grandmaster Clock.

3.1.19 holdover or holdover mode: A clock A, previously synchronized/syntonized to another clock B (normally a primary reference or a Master Clock) but whose frequency is determined in part using data acquired while it was synchronized/syntonized to B, is said to be in holdover or in the holdover mode as long as it is within its accuracy requirements.

3.1.20 implemented message timestamp point: A point within a Precision Time Protocol event message serving as a timestamp-capture point in the message. A timestamp is captured at the instant an implemented message timestamp point passes a point possibly removed from the reference plane. Depending on the implementation, the implemented message timestamp point is or is not equivalent to the message timestamp point. See: **message timestamp point**.

NOTE—The relationship between the two is specified in 7.3.4.2.

3.1.21 integrity check value (ICV): The result of a cryptographic function used to ensure that unauthorized modifications of a message are detected (e.g., the value might be the result of the keyed hash function HMAC-SHA256).

NOTE—See FIPS PUB 180-4:2015.

3.1.22 key: In the context of a security association, refers to the secret used in the cryptographic calculation of the integrity check value (ICV).

3.1.23 key distribution center (KDC): The functionality of a (typically) centralized authority to which peers authenticate before distribution of security parameters. For instance, KDCs are used to distribute security parameters for group-based communication.

3.1.24 L1 rx clock signal: A clock signal that is recovered from the reception of data from the medium.

3.1.25 L1 tx clock signal: A clock signal that is used in the transmission of data over the medium.

3.1.26 Local Clock: A physical clock that is used in the operation of the PTP Protocol.

- The Local PTP Clock may be derived from or identical to the Local Clock;
- The Local Clock may be syntonized to a reference clock signal.

NOTE—See 12.2.

3.1.27 Local Clock signal: The clock signal of a Local Clock.

3.1.28 Local PTP Clock: The clock of a PTP Instance that provides the local estimate of the time of its Grandmaster Clock, that is it is synchronized to the time of the Grandmaster Clock. It is either a physical or a mathematical clock, and it provides PTP or ARB (arbitrary) time.

NOTE—See 7.2.1.

3.1.29 Local PTP Clock signal: The clock signal of a Local PTP Clock.

3.1.30 management mechanism: A management protocol used for configuring and/or monitoring PTP Nodes and PTP Instances.

NOTE—This term refers to the PTP management messages of this standard, as well as to the management protocols defined in other standards (e.g., SNMP [B27]⁹ and NETCONF [B20]).

3.1.31 Master Clock: In the context of a single PTP Communication Path, the Local PTP Clock of an Ordinary Clock or Boundary Clock that is the source of time to which all other Local PTP Clocks on that PTP Communication Path synchronize.

3.1.32 Master PTP Instance: A PTP Instance containing a Master Clock.

NOTE—A Master PTP Instance propagates the traceability information from the Grandmaster PTP Instance throughout the associated PTP Communication Path.

3.1.33 message timestamp point: A point within a PTP event message serving as a reference point in the message. A timestamp is defined by the instant a message timestamp point passes the reference plane of a PTP Instance.

3.1.34 mixed-step PTP Instance: A PTP Instance where some PTP Ports can be one-step PTP Ports and one or more PTP Ports can be two-step PTP Ports, see one-step PTP Port and two-step PTP Port.

3.1.35 multicast communication: A communication model in which each PTP message sent from any PTP Port is received and processed by all PTP Ports on the same PTP Communication Path.

⁹ The numbers in brackets correspond to those of the bibliography in Annex Q.

3.1.36 Mutable field: A field that can be modified on retransmission by PTP Instances. Handling of mutable fields in the context of security depends on the key management system.

3.1.37 one-step port semantics: Semantics where a transmitted Sync or Pdelay_Resp message will not be followed by a Follow_Up or Pdelay_Resp_Follow_Up message, respectively.

3.1.38 one-step PTP Instance: A PTP Instance where all PTP Ports are one-step PTP Ports.

3.1.39 one-step PTP Port: A PTP Port that obeys one-step port semantics.

3.1.40 Ordinary Clock: A PTP Instance that has a single PTP Port in its domain and maintains the timescale used in the domain. An Ordinary Clock can serve as a source of time, that is, contain a Master Clock, or alternatively, the Local PTP Clock of an Ordinary Clock can be synchronized, that is, be a Slave Clock, to the Local PTP Clock of a Boundary Clock or another Ordinary Clock in the domain.

3.1.41 Parent PTP Instance: Within a domain, an Ordinary Clock or Boundary Clock A, is the Parent PTP Instance of another Ordinary Clock or Boundary Clock B, if the Local PTP Clock of A is the Master Clock of the Slave Clock in B.

3.1.42 peer-to-peer port: A PTP Port that is configured to use the peer-to-peer delay mechanism.

3.1.43 peer-to-peer Transparent Clock: A Transparent Clock that, in addition to providing PTP event transit time information, also corrects for the propagation delay of the PTP Link connected to the PTP Port receiving the Sync message. In the presence of peer-to-peer Transparent Clocks, delay measurements between Slave Clocks and the Master Clock are performed using the peer-to-peer delay mechanism.

3.1.44 phase change rate: The observed rate of change in the measured time with respect to the reference time. The phase change rate is equal to the frequency offset between the measured frequency and the reference frequency. *See also:* **frequency offset**

3.1.45 phase offset: Let $T_A(t)$ and $T_B(t)$ be the times indicated by clocks associated with a clock signal A and a clock signal B, respectively, as a function of time t . Then the phase offset of A with respect to B is $T_A(t) - T_B(t)$. A positive value of the phase offset means that A is in advance of B. Phase offset can exist between clocks, clock signals, as well as a clock and a clock signal.

NOTE—Evaluation of the phase offset does not require an existence of clocks associated with the respective clock signals. A phase offset can be directly evaluated by measuring the time interval between the relevant periodic events of clock signals A and B. Given the definition, in such a measurement, a positive value of the phase offset means that the periodic event of the clock signal A occurs earlier in time than the corresponding event of the clock signal B.

3.1.46 portNumber: An index identifying a specific PTP Port on a PTP Instance.

3.1.47 precision: See: **accuracy**.

3.1.48 Precision Time Protocol (PTP): The protocol defined by IEEE Std 1588™. As an adjective, it indicates that the modified noun is specified in or interpreted in the context of IEEE Std 1588™.

3.1.49 primary reference: A source of time and/or frequency that is traceable to international standard(s). *See also:* **traceability**.

3.1.50 Profile (or PTP Profile): A document, or a portion of a document, specifying the set of PTP features and attribute values applicable to a PTP Instance, and written by an organization following the specifications of 20.3 of IEEE Std 1588™-2019.

3.1.51 PTP Communication: Information used in the operation of the PTP protocol, transmitted in a PTP message.

3.1.52 PTP Communication Path: Within a domain, the signaling path portion of a particular PTP Network connecting Ordinary Clocks and Boundary Clocks with no intervening Boundary Clocks. However, there can be Transparent Clocks or non-PTP devices in a PTP Communication Path.

NOTE—See Figure 2.




Figure 2—PTP Communication Path

3.1.53 PTP Instance: An instance of the PTP protocol, operating in a single device, within exactly one domain. A PTP Instance implements those portions of the standard indicated as applicable to an Ordinary Clock, a Boundary Clock, or a Transparent Clock.

3.1.54 PTP Instance Time: Time measured on the timescale maintained by the Local PTP Clock of the PTP Instance. The timescale of a Local PTP Clock is either the timescale PTP or the timescale ARB. See 7.2.1 of IEEE Std 1588™-2019.

NOTE—PTP Instance Time is not available in a Transparent Clock that does not maintain a Local PTP Clock.

3.1.55 PTP Link: Within a domain, a network segment between two PTP Ports using the peer-to-peer delay mechanism of this standard. The peer-to-peer delay mechanism is designed to measure the propagation time over such a link.

NOTE 1—In many applications, this network segment can be a direct physical link between the PTP Nodes including the above PTP Ports.

NOTE 2—A PTP Link between PTP Ports of Boundary Clocks or Ordinary Clocks is also a PTP Communication Path, see PTP Communication Path.

3.1.56 PTP management message: A PTP message defined in IEEE Std 1588™ for the purpose of configuring and/or monitoring PTP Nodes and PTP Instances.

NOTE—This term does not refer to messages of management protocols defined in other standards (e.g., SNMP and NETCONF).

3.1.57 PTP Management Node: A device that configures and/or monitors Boundary Clocks, Ordinary Clocks and/or Transparent Clocks in one or more domains using the PTP management messages of IEEE Std 1588.

3.1.58 PTP message: One of the messages defined in IEEE Std 1588™.

NOTE—See 6.4.

3.1.59 PTP Network: A network consisting of a combination of PTP Nodes and possibly non-PTP devices and/or PTP Management Node(s). Non-PTP devices include, for example, some bridges, routers, and other infrastructure devices, and possibly devices such as computers, printers, and other application devices.

3.1.60 PTP Node: A device that contains one or more PTP Instances and/or PTP services (e.g., Common Mean Link Delay Service).

3.1.61 PTP Port: A logical access point of a PTP Instance for PTP Communication to the communications network.

3.1.62 recognized standard time source: A source external to Precision Time Protocol (PTP) that provides time and/or frequency as appropriate, and that is traceable to the international standards laboratories maintaining clocks that form the basis for the International Atomic Time (TAI) and Coordinated Universal Time (UTC) timescales. Examples of these are GNSS systems (e.g. GPS, GLONASS), and national time laboratories (e.g. National Institute of Standards and Technology in the US, National Physical Laboratory in the UK).

3.1.63 requester: In the context of the peer-to-peer delay mechanism, the PTP Port implementing the peer-to-peer delay mechanism that initiates the operation of the mechanism by sending a Pdelay_Req message.

3.1.64 responder: In the context of the peer-to-peer delay mechanism, the PTP Port responding to the receipt of a Pdelay_Req message as part of the operation of the peer-to-peer delay mechanism.

3.1.65 retransmit: In the context of a PTP Instance processing PTP messages (e.g. the PTP stack running above the MAC layer), retransmission of a PTP message means that the PTP-defined contents of a received PTP message, possibly updated, is the PTP-defined content of new PTP message(s) transmitted on one or more PTP Ports distinct from the PTP Port on which the original PTP message was received.

3.1.66 Security Association (SA): Description of the set of security parameters necessary to provide security services (e.g., authentication and integrity) between different entities sharing the same SA. The security parameters of the SA are typically negotiated or signaled in the context of the key management protocol.

3.1.67 Security Association Database (SAD): Database containing the sets of SAs with their associated security parameters. The content of the database is dynamic in the context of the application being secured.

3.1.68 Security Parameter Pointer (SPP): Identifier of a specific SA. If carried in the application protocols it is used to point to the associated SA in the SAD and the SPD.

3.1.69 Security Policy Database (SPD): Database containing the required security services (e.g., message integrity required for all messages) with their specifics (e.g., at least support for HMAC-SHA256 and an associated key for ICV calculation) for an application. The SPD is typically involved in the establishment of an SA to ensure the adherence to a given policy. The SPD is configured and content is static between any reconfiguration.

3.1.70 Special Port: A Special Port is a PTP Port that implements a mechanism for transferring time over a network where the time transfer is not based on the use of PTP timing messages.

NOTE—See 11.5.1.

3.1.71 Slave Clock: In the context of a single PTP Communication Path, the Local PTP Clock of an Ordinary Clock or Boundary Clock that synchronizes to the Local PTP Clock of the Master PTP Instance on that PTP Communication Path.

3.1.72 Slave PTP Instance: A PTP Instance containing a Slave Clock.

NOTE—A Boundary Clock can be both a Master PTP Instance and a Slave PTP Instance. See **Master PTP Instance**, **Slave PTP Instance** and **Boundary Clock**.

3.1.73 specified-by-design: An attribute, field, or state variable with a value inherent in the implementation or specified by IEEE Std 1588™ or the applicable PTP Profile as opposed to the case where the value is subject to change by configuration or the operation of the protocol.

3.1.74 stability (of a clock or clock signal): A measure of the variations over time of the frequency error (of the clock or clock signal). The frequency error typically varies with time due to aging and various environmental effects, for example, temperature.

3.1.75 synchronized clocks: Absent relativistic effects, two clocks are synchronized to a specified uncertainty if they have the same epoch and their measurements of the time of the same single event occurring at an arbitrary instant differ by no more than that uncertainty.

3.1.76 syntonized clocks: Absent relativistic effects, two clocks are syntonized to a specified uncertainty if the duration of the second is the same on both, which means the time as measured by each advances at the same rate within the specified uncertainty. The two clocks might or might not share the same epoch.

3.1.77 time counter: A digital time representation of the time that is incremented by the period of a clock signal at each periodic event of the clock signal. *See:* **clock signal**.

3.1.78 timeout: A mechanism for terminating requested activity that, at least from the perspective of the entity requesting termination, does not complete within the specified time.

3.1.79 timescale: A measure of elapsed time since an epoch.

3.1.80 Timestamping Clock: The Timestamping Clock is used in generating ingress and egress timestamps for PTP event messages and is either the Local PTP Clock or the Local Clock.

NOTE—See: 7.3.4.3.

3.1.81 traceability: A property of the result of a measurement or the value of a standard whereby it can be related to stated references, usually national or international standards, through an unbroken chain of comparisons all having stated uncertainties. (Adapted from the International Vocabulary of Basic and General Terms in Metrology Second Edition 1993 Definition 6.10). For purposes of this standard, the term “unbroken chain of comparisons” means that if measurements were performed comparing the difference between the stated reference and any point in the chain (for example, the time difference between time in the GNSS system and the time in a Grandmaster Clock), the results of these measurements would be within the stated uncertainties. For example, as long as a measurement, if performed, were known to be within the stated uncertainty, it does not matter whether the Grandmaster PTP Instance is actively disciplining its Local PTP Clock to the GNSS system or whether the Local PTP Clock is in holdover.

3.1.82 traceable: The time or frequency of a clock if measurements of these quantities have the traceability property.

3.1.83 translation device: A Boundary Clock or, in some cases, a Transparent Clock that translates the PTP messages between regions implementing different transport mappings, between different versions of this standard, or different Precision Time Protocol (PTP) Profiles.

3.1.84 Transparent Clock: A PTP Instance that measures the time for a PTP event message to transit the PTP Instance and provides this information to PTP Instances receiving this PTP event message. Peer-to-peer Transparent Clocks also correct for PTP Link delay. *See also:* **end-to-end Transparent Clock**, **peer-to-peer Transparent Clock**

3.1.85 two-step port semantics: Semantics where a transmitted Sync or Pdelay_Resp message will be followed by a Follow_Up or Pdelay_Resp_Follow_Up message, respectively.

3.1.86 two-step PTP Instance: A PTP Instance where all PTP Ports are two-step PTP Ports.

3.1.87 two-step PTP Port: A PTP Port that obeys two-step port semantics.

3.2 Acronyms and abbreviations

ARB	arbitrary
BMCA	best master clock algorithm
CID	Company identification (allocated by the IEEE; see 7.5.2.2.1)
CMLDS	Common Mean Link Delay Service
CP	Communication Profile [according to IEC 61784-1:2007]
CPF	Communication Profile Family [according to IEC 61784-1:2007]
E2E	end-to-end
EEC	synchronous Ethernet Equipment Clock
EEC1	EEC Option 1
EEC2	EEC Option 2
EPON	Ethernet Passive Optical Network
GDOI	Group Domain of Interpretation
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
IANA	Internet Assigned Numbers Authority
ID	identification
IEC	International Electrotechnical Commission
IEEE RA	IEEE Registration Authority
IERS	International Earth Rotation and Reference Systems Service
IETF	Internet Engineering Task Force
IPv4, IPv6	Internet Protocol version 4 / 6
ISO	International Organization for Standardization
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
L1	layer-1
MAC	media access control [according to IEEE Std 802.3]
MDMI	Media-Dependent, Media-Independent
MJD	Modified Julian Day
MTIE	Maximum Time Interval Error
MUC	Medium Under Calibration
NIST	National Institute of Standards and Technology (see www.nist.gov)
NTP	Network Time Protocol (see IETF RFC 5905 [B19])
NUI	Network Unique Identifier
OID	Object Identifier
OUI	Organizationally Unique Identifier (allocated by the IEEE; see 7.5.2.2.1)
P2P	peer-to-peer
PDU	Protocol Data Unit
PHY	physical layer [according to IEEE Std 802.3]
PIUC	PTP Instance Under Calibration
PLL	Phase Locked Loop
POSIX	Portable Operating System Interface (see ISO/IEC 9945 [B30])
ppm	parts per million
PPS	pulse per second
PTP	Precision Time Protocol
QSDO	Qualified Standards Development Organization
RFC	Request For Comments
rx	receive
SFP	small form-factor pluggable
TAI	International Atomic Time
TC	traffic class
TDEV	Time deviation
TESLA	Time Efficient Stream Loss-Tolerant Authentication
TLV	type, length, value
tx	transmit
UCMM	UnConnect Message Manager

UDP/IP	User Datagram Protocol [see IETF RFC 768 (1980)]/Internet Protocol [see IETF RFC 791 (1981)]
UTC	Coordinated Universal Time



4. Conventions

4.1 Descriptive lexical form syntax

4.1.1 Lexical form syntax

A lexical form refers to:

- a) A name,
- b) A data type.

The conventions illustrated in the following list regarding lexical forms are used in this standard:

- **Type names:** For example, ClockQuality (no word separation, initial letter of each word capitalized).
- **Enumeration members and global constants:** For example, ATOMIC_CLOCK (underscore word separation, all letters capitalized).
- **Fields within PTP messages, instances of structures, and variables:** For example, secondsField, clockQuality, clockIdentity (two-word field names at a minimum, no word separation, initial word not capitalized, initial letter capitalization on subsequent words).
- **Members of a structure:** For example, clockQuality.clockClass (structure name followed by a period followed by the member name).
- **Data set name:** For example, defaultDS, parentDS, portDS, currentDS, timePropertiesDS (no word separation, initial letter of each word not capitalized followed by the letters DS).
- **Data set members:** For example, defaultDS.clockQuality.clockClass (data set name followed by a period followed by a member name followed by a period followed by a member name) or defaultDS.instanceEnable (data set name followed by a period followed by a member name).
- **PTP message names:** For example, Sync, Delay_Req (underscore word separation, initial letter of each word capitalized).
- **<localNameForSomething>:** Text enclosed in angle brackets < > is used where the standard needs to refer to something whose syntax or lexical form is dependent on the local implementation and language, for example, in 5.3.1 in defining data types. It is also used to express an abstract quantity as part of a description, for example, <delayAsymmetry> in 7.4.2 in connection with Figure 27, or as a computed quantity used in the operation of the protocol, for example, <delayAsymmetry> in 8.2.15.4.8 in defining the value of portDS.delayAsymmetry, or in 10.2.2.1.1 for computing the value of the correctionField of a Sync message. The <localNameForSomething> is also used when discussing artifacts, for example, measurements or computations, in the abstract as opposed to normative specifications. Angle bracket notation is also used in defining a type, to indicate what is filled in. For example: struct <StructureName> (see 5.3.1).
- **Names of entities and subsystems in the model architecture:** For example, Boundary Clock, PTP Node, Clock Source, PTP Common Core (separate words all with initial Capitalization).

When a lexical form appears in text, as opposed to in a type, or a format definition, the form is to be interpreted as singular, plural, or possessive as appropriate to the context of the text.

4.2 Word usage

4.2.1 Shall

The word “shall,” which is equivalent to “is required to,” is used to indicate mandatory requirements, strictly to be followed to conform to the standard and from which no deviation is permitted.

4.2.2 Must

The word “must” indicates an unavoidable situation.

4.2.3 Should

The word “should,” which is equivalent to “is recommended that,” is used to indicate the following:

- Among several possibilities, one is recommended as particularly suitable, without mentioning or excluding others.
- That a certain course of action is preferred but not required.
- That (in the negative form) a certain course of action is deprecated but not prohibited.

4.2.4 May

The word “may,” which is equivalent to “is permitted,” is used to indicate a course of action permissible within the limits of the standard.

4.2.5 Can

The word “can,” which is equivalent to “is able to,” is used to indicate possibility and capability, whether material or physical.

4.2.6 Optional

Clauses, annexes, and text marked optional are not required to be implemented. If the option is implemented, then all specifications of the clause, annex, or text, respectively, shall be implemented according to this standard.

NOTE—This definition is recursive, which means that options within options obey these same rules.

4.2.7 Reserved

The word “reserved” indicates:

- If used in an assignment of values to an enumeration or an attribute that the values indicated are reserved for use in future editions of this standard and shall not be used for any other purpose.
- If used in a field of a PTP message that the field is reserved for use in future editions of this standard. The field shall be present in the PTP message with the size specified. No interpretation of reserved fields is to be made for this edition of this standard, and the fields shall not be used for any other purpose.

4.2.8 Deprecated

The term “deprecated” indicates that a feature of the standard is permitted to be used in compliant devices, but it is no longer recommended or required. The function realized by the deprecated feature can now be realized using an alternate, superior option. Deprecated features may serve an important role in maintaining backward compatibility.

4.2.9 Experimental

The word “experimental” indicates that:

- Any artifact designated in this standard as experimental is intended for use in conjunction with protocol development within a privately administered development network, for example, within a network that has no wide area connectivity. Within that network, a local administrator is free to use artifacts marked as experimental for protocol development purposes.
- These experimental artifacts shall not be used in protocols or products that are to be released for use in the wider networking community, as freeware, shareware, or any part of a company’s commercial product offering, and shall not be specified in a standard or be permanently assigned for use with a given protocol or protocols.
- Prior to deployment in an environment outside the developing organization’s administrative control or for public release as a standard or product, the experimental artifact and any specifications dependent on the artifact shall either:
 - Be submitted to and approved by the IEEE 1588 Working Group¹⁰ as a part of the IEEE 1588 standard. In this case, the specification of the artifact and any dependent specifications shall be a part of IEEE Std 1588 and not of any other standard, or
 - Be specified in a PTP Profile subject to the restrictions of 20.3.1.2, for example, as an alternate BMCA (see 9.3.1), or organization-specific TLV (see 14.3)
- Implementers must be aware that experimental artifacts, for example, experimental TLV values (see Table 52 in 14.1.1) or experimental sdoId values (see Table 2 in 7.1.4), can be used by anyone subject to the conditions of this clause. The use of such artifacts in violation of 4.2.9 can lead to interoperability issues in PTP Networks.

4.3 Behavioral specification notation

State transition diagrams are used to specify behavioral characteristics as illustrated in Figure 3. Each state transition diagram is composed of the following components:

- Named boxes, representing states
- Directed arrows, indicating transitions from one state to the next

Each transition is labeled with:

- a) Enabling event or predicate label for a transition
- b) Transition action label for a transition

¹⁰ IEEE Precise Networked Clock Synchronization Working Group. Also referred to in this document as the IEEE 1588 Working Group (<https://standards.ieee.org/project/1588.html>).




Figure 3—Mealy state transition diagram

The notation used describes state transition diagrams using the Mealy style, where actions are associated with the transition from one state into another.

Events, for example, “event_1,” “event_2,” and “event_3,” identify the inputs to the state machine. They can be operation requests and responses or internal occurrences such as timer expirations.

Predicates, for example, “event_1 OR event_2,” identify enabling conditions for transitions. The first predicate encountered, evaluated from left to right, which is TRUE, selects the transition to execute and therefore the next state.

Transition actions, for example, “result_1,” are the actions that are executed before transitioning to the next state.

The next state identifies the state for the state machine after the selected transition action completes. The value of the current state changes as the transition to the next state occurs.

A bold line for a state box indicates that the box represents multiple states. Any transition shown that begins and terminates in such a state box indicates that there has been no change in state.

Transitions, for example, the transition resulting in result_3, which have no indicated enabling conditions, occur via unspecified mechanisms. Unless otherwise stated in this standard, the events giving rise to these mechanisms are implementation specific and outside the scope of the standard.

A transition into a state machine, for example, “(i),” is indicated by a transition arrow that has no source state. A transition out of a state machine, for example, “(d),” is indicated by a transition arrow with no destination state.

NOTE—For example, as a result of either event_1 or event_2 becoming TRUE, State_1 is replaced with the value of the next state. In this example, the next state is State_2, which is specified as the name of the state box that is the target of the transition arrow. Before the transition, result_1 occurs. “event_3” can occur in either State_1 or State_2. The state is unchanged, but an action, result_2, occurs as the result of event_3.

5. Data types and on-the-wire formats

5.1 General

The data types specified for the various PTP variables and PTP message fields define logical properties that are necessary for correct operation of the protocol or interpretation of PTP message content.

Implementations are free to use any internal representation of PTP data types; however, the internal representation shall not change the semantics of any quantity visible via PTP Communications or the semantics of any specified operation of the protocol.

5.2 Primitive data type specifications

All nonprimitive PTP data types are derived from the primitive types listed in Table 1. These types are not tied to any specific programming language. The essential properties of each type shall be as follows:

- **Integer:** All integer types are of finite length as indicated by the number associated with each, for example, UInteger48, and as signed or unsigned as indicated by the absence or presence of a leading U. Numbers with these data types obey the laws of arithmetic within the range represented by the length. Arithmetic operations are treated as modulo the capacity of the data type; for example, the sum of two UInteger48 values is computed modulo 2^{+48} . Signed integers are represented in two's complement form.
- **Enumeration:** All enumerations are of finite field length as indicated by the number associated with each, for example, Enumeration4. Unless otherwise stated in this standard, the only interpretations of the bit pattern in the enumeration field are the associations between the bit patterns and the assigned meanings of the enumeration.
- **Boolean:** The only interpretation is as logical values within the context of Boolean algebra.
- **Nibble and Octet:** These are 4- and 8-bit fields, respectively. The only interpretations are those explicitly defined within this standard.
- **Float:** IEEE 754™ floating-point number, with a length that specifies the format. Support for floating-point data types is optional because usage of those data types is limited to the specifications of optional features.



Table 1—Primitive PTP data types

Data type	Definition
Boolean	TRUE or FALSE
Enumeration4	4-bit enumerated value
Enumeration8	8-bit enumerated value
Enumeration16	16-bit enumerated value
UInteger4	4-bit unsigned integer
Integer8	8-bit signed integer
UInteger8	8-bit unsigned integer
UInteger12	12-bit unsigned integer
Integer16	16-bit signed integer
UInteger16	16-bit unsigned integer
Integer32	32-bit signed integer
UInteger32	32-bit unsigned integer
UInteger48	48-bit unsigned integer
Integer64	64-bit signed integer
Nibble	4-bit field not interpreted as a number
Octet	8-bit field not interpreted as a number
Float64	IEEE 754 binary64 (64-bit double-precision floating-point format)

5.3 Derived data type specifications

5.3.1 General

Arrays of any of the primitive data types are represented in the format <data type>[lengthField] <label>, where [lengthField] indicates the number of instances of the data type in the array and <label> is the lexical name for the array data type so defined.

Structures consisting of an ordered list of members are indicated with the following syntax:

```
Struct <StructureName>
{
  <DataType1> <memberName1>;
  <DataType2> <memberName2>;
  ...
};
```

where <StructureName> is the lexical name for the data type so defined, <DataType1> is the data type of the first member, <memberName1> is the lexical name of the first member, and so forth.

The syntax Typedef <DataType> <TypeName> is interpreted as defining a derived data type with the same properties as the data type defined by <DataType> but with a new name given by <TypeName>.

The syntax Typedef <DataType> [lengthField] <TypeName> is interpreted as defining a derived data type consisting of an array of elements of type <DataType> but with a new name given by <TypeName>.

For the derived data types, the specifications of the semantics and values of all members of structs and arrays are found in the clauses specifying the use of the struct or array. This subclause specifies only the data types.

5.3.2 TimeInterval

The TimeInterval type represents time intervals.

```
Typedef Integer64 TimeInterval;
```

TimeInterval is the time interval expressed in nanoseconds, multiplied by 2^{+16} .

Positive or negative time intervals outside the maximum range of this data type shall be encoded as the largest positive and negative values of the data type, respectively.

For example, 2.5 ns is expressed as 0000 0000 0002 8000₁₆.

NOTE—In the 2008 edition of this standard, TimeInterval was defined as a Struct. There are no backward compatibility issues because the on-the-wire format is the same.

5.3.3 Timestamp

The Timestamp type represents a positive time with respect to the epoch.

```
Struct Timestamp
{
  UInteger48 secondsField;
  UInteger32 nanosecondsField;
};
```

The secondsField member is the integer portion of the timestamp in units of seconds.

The nanosecondsField member is the fractional portion of the timestamp in units of nanoseconds.

The nanosecondsField member is always less than 10^9 .

For example:

+2.000000001 seconds is represented by secondsField = 0000 0000 0002₁₆ and nanosecondsField = 0000 0001₁₆.

5.3.4 ClockIdentity

The ClockIdentity type identifies unique entities within a PTP Network, for example, a PTP Instance or an entity of a common service.

```
Typedef Octet[8] ClockIdentity;
```

5.3.5 PortIdentity

The PortIdentity type identifies a PTP Port (see 7.5) or a Link Port (see 16.6.1).

```
Struct PortIdentity
{
  ClockIdentity clockIdentity;
  UInteger16 portNumber;
};
```

5.3.6 PortAddress

The PortAddress type represents the protocol address of a PTP Port.

```
Struct PortAddress
{
  Enumeration16 networkProtocol;
  UInteger16 addressLength;
  Octet[addressLength] addressField;
};
```

The value of the networkProtocol member shall be taken from the networkProtocol enumeration (see 7.4.1).

The addressLength is the length in octets of the address. The range shall be 1 to 16 octets.

The addressField member holds the protocol address of a PTP Port in the format defined by the mapping annex of the protocol as identified by the networkProtocol member. The most significant octet of the addressField is mapped into the octet of the addressField member with index 0.

5.3.7 ClockQuality

The ClockQuality represents the quality of a clock.

```
Struct ClockQuality
{
  UInteger8 clockClass;
  Enumeration8 clockAccuracy;
  UInteger16 offsetScaledLogVariance;
};
```

5.3.8 TLV

The TLV type represents TLV extension fields.

```
Struct TLV
{
  Enumeration16 tlvType;
  UInteger16 lengthField;
  Octet [lengthField] valueField;
};
```

The length of all TLVs shall be an even number of octets.

5.3.9 PTPText

The PTPText data type is used to represent textual material in PTP messages.

```
Struct PTPText
{
  UInteger8 lengthField;
  Octet[lengthField] textField;
};
```

The textField member shall be encoded as UTF-8 symbols as specified by ISO/IEC 10646:2003. The most significant octet of the leading text symbol shall be the element of the array with index 0.

NOTE—A single UTF-8 symbol can be 1 to 4 octets long. Therefore, the lengthField value can be larger than the number of symbols. A symbol is analogous to a character.

PTPText is also used to represent textual strings in management mechanisms that do not use PTP messages. When used in this context, PTPText shall represent the string data type of the corresponding management standards (not the preceding struct).

5.3.10 FaultRecord

The FaultRecord type is used to construct fault logs.

```
Struct FaultRecord
{
  UInteger16 faultRecordLength;
  Timestamp faultTime;
  Enumeration8 severityCode;
  PTPText faultName;
  PTPText faultValue;
  PTPText faultDescription;
};
```

The faultRecordLength member shall indicate the number of octets in the FaultRecord not including the 2 octets of the faultRecordLength member.

5.3.11 RelativeDifference

The RelativeDifference type represents the relative difference between two numeric values.

```
Typedef Integer64 RelativeDifference;
```

RelativeDifference is the relative difference expressed as a dimensionless fraction and multiplied by 2^{+62} , with any remaining fractional part truncated.

NOTE 1—For example, 0.00024 is expressed as 0003 EEA2 09AA A3AD₁₆. The maximum range of RelativeDifference is (-2^{+1}) to $(2^{+1}-2^{-62})$, which is sufficient to store a fractional number.

NOTE 2—When using values with the data type RelativeDifference, in calculations in which other values are expressed in units of 2^{-16} ns (data type TimeInterval) or nanoseconds, the values need to be appropriately scaled before performing the calculations. In some implementations, it might be desired to avoid floating point arithmetic when working with expressions that involve the two data types, that is, RelativeDifference and TimeInterval. This can be accomplished by appropriately scaling the input raw values (expressed in dimensionless units of 2^{+62} , units of 2^{-16} ns, or units of nanoseconds) or the product of the computation. In the latter case, the raw values can be used if the result of the computation is appropriately scaled. For example, if a quantity of data type RelativeDifference is multiplied by a quantity with units of nanoseconds, and the result is desired in units of 2^{-16} ns, the raw values can be multiplied and then the product can be divided by 2^{+46} to obtain the result in units of 2^{-16} ns.

5.3.12 CommunicationCapabilities

The CommunicationCapabilities structure lists the communication modes in which a PTP Port can operate.

```
Struct CommunicationCapabilities
{
  Boolean multicastCapable;
  Boolean unicastCapable;
  Boolean unicastNegotiationCapable;
  Boolean unicastNegotiationRequired;
};
```

5.4 On-the-wire formats

5.4.1 General

PTP protocol data units consist of the PTP messages defined in this standard. The internal ordering of the fields of the PTP protocol data units is specified in 5.4.2 to 5.4.4.

5.4.2 Primitive data types

Numeric primitive data types defined in 5.2 shall be formatted with the most significant octet nearest the beginning of the protocol data unit followed in order by octets of decreasing significance.

The Boolean data type TRUE shall be formatted as a single bit equal to 1 and FALSE as a single bit equal to 0.

Enumerations of whatever length shall be formatted as though the assigned values are unsigned integers of the same length; for example, Enumeration16 shall be formatted as though the value had type UInteger16.

5.4.3 Arrays of primitive types

All arrays shall be formatted with the member having the lowest numerical index closest to the start of the protocol data unit followed by successively higher numbered members. In octet arrays, the octet with the lowest numerical index is termed the most significant octet.

When a field containing more than one octet is used to represent a numeric value, the most significant octet shall be nearest the start of the protocol data unit, followed by successively less significant octets.

When a single octet contains multiple fields of primitive data types, the bit positions within the octet of each of the primitive types as defined in the PTP message field specification shall be preserved. For example, the first field of the header of the PTP messages is a single octet composed of two fields, one of type Nibble bits 4 to 7, and one of type Enumeration4 bits 0 to 3 (see 13.3.1).

5.4.4 Derived data types

Derived data types defined as structs shall be formatted with the first member of the struct closest to the beginning of the protocol data unit followed by each succeeding member. Each member shall be formatted according to its data type.

Derived data types defined as typedefs shall be formatted according to its referenced data type.

5.4.5 Mapping of PTP protocol data units into their on-the-wire formats

Unless otherwise specified, PTP protocol data units shall be mapped to and from their on-the-wire format based on the rules of the underlying physical layer transport. Any exceptions are specified in one of the transport-specific annexes of this standard or in the applicable PTP Profile.

NOTE—PTP protocol mechanisms operate in the upper layers of the protocol stack (i.e., PTP is an “application” that uses the services of the network or link layer). The physical layer transport dictates the on-the-wire format.

6. Clock synchronization model

6.1 General requirements on implementations

6.1.1 General

This clause provides a model for understanding the operation of the Precision Time Protocol. The exact specifications for this operation are found in subsequent clauses.

This standard specifies a detailed model of the Precision Time Protocol. Implementers of PTP Instances may use internal representations, organizations, and procedures other than those presented in this document. The externally visible behavior of implementations of PTP Instances shall, in all respects, be identical to the behavior, including the externally visible results of computations, timing considerations, state evolutions, PTP message formats, PTP message contents, and data types, as specified by the models and specifications of this standard.

This standard provides the following classes of alternatives and options:

- Alternative transport mechanisms, of which at least one must be chosen (e.g., Annex C, Annex D, and Annex E)
- Alternative specifications for implementing normative requirements, for example, flexibility in conveying information in blade architectures, and the choice between methods of determining message propagation time
- Optional features that aid in constructing applications, for example, security or high accuracy
- Optional features that fine tune or enhance performance of the best master clock algorithm, for example, unicast discovery

In all cases, which of the alternatives is applicable, or which of the options is implemented and/or enabled/disabled, shall be specified by one or more of the following:

- The manufacturer
- The applicable PTP Profile
- One of the optional features specifying that another option or alternative be used

If an option or alternative is implemented, any associated optional data sets shall be implemented irrespective of whether the option or alternative is enabled or disabled. In this context, “implemented” means that a PTP Instance contains the executable, that is, the code, data sets, hardware, or whatever is required to execute the option or alternative. In the absence of an enable/disable feature and if the option is implemented, the executable is present and functioning, that is, the option is enabled by design.

The internal representations, organization, and procedures utilized by implementations when performing arithmetic operations (e.g., addition and subtraction) on numerical attributes (e.g., Integer data type) need only suffice to achieve the level of precision required to meet their designated performance levels. Designated performance levels can be imposed by a requirement within this standard (e.g., within procedures involved in determining the designated values of specified numerical attributes), and they may be imposed by a PTP Profile or the needs of the target application.

When implementing arithmetic operations, implementers must observe and correct for differences in data type, units, and in some cases scaling, for example, differences between data types of timestamps and correction fields.

Management configuration of numerical attributes that will be later used in arithmetic operations may be limited to a dynamic range smaller than the dynamic range enabled by the data type defined by the management interface of the attribute. The effective configurable dynamic range need only suffice to achieve the designated performance levels. Designated performance levels can be imposed by a requirement within this standard (e.g., within procedures involved in determining the designated values of specified numerical attributes), and they may be imposed by a PTP Profile or the needs of the target application.

6.1.2 Applicability and contents of the standard

The PTP standard specifies a clock synchronization protocol. This protocol is applicable to distributed systems consisting of two or more PTP Instances, communicating over a PTP Network. PTP Instances are modeled as containing a real-time clock that may be used by applications associated with the PTP Instance for various purposes, such as generating timestamps for data or ordering events managed by a PTP Instance. The protocol provides a mechanism for synchronizing the clocks of participating PTP Instances to a high degree of accuracy and precision. A PTP Node can contain multiple PTP Instances.

This standard specifies:

- a) The Precision Time Protocol
- b) The PTP Node, PTP Network, PTP Instance, and communication properties necessary to support PTP

6.1.3 Names for major entities defined by the standard

Numerous terms are defined in Clause 3 that refer to major entities specified by the standard. These are:

- a) PTP Instance
- b) PTP Node
- c) Boundary Clock
- d) Ordinary Clock
- e) Transparent Clock
- f) Local Clock, Local PTP Clock
- g) PTP Management Node

The relationships between these entities are illustrated in Figure 4.




Figure 4 — Relationship of PTP entities

From Figure 4, each PTP Node can contain PTP Instances, that is, Boundary Clocks, Ordinary Clocks, or Transparent Clocks. For a single domain per 7.1, each such PTP Instance implements the functionality specified for the specific PTP Instance, that is, a Boundary Clock, Ordinary Clock, or Transparent Clock. For all PTP Instances, this functionality includes the PTP protocol stack specified for the type of the PTP Instance.

PTP Node “A” of Figure 4 illustrates PTP Instances, each of which contains either both a Local PTP Clock and a Local Clock or else only a Local Clock. The exact configuration depends on application requirements. For example, an Ordinary Clock will contain a Local PTP Clock and a Local Clock. By contrast, a Transparent Clock only needs a Local Clock although a Local PTP Clock can be present.

PTP Node “B” of Figure 4 illustrates PTP Instances, each of which contains a Local PTP Clock and/or access to a Local Clock common to all PTP Instances of the PTP Node. The exact configuration depends on application requirements. For example, an Ordinary Clock will contain a Local PTP Clock and have access to the Local Clock. By contrast, a Transparent Clock only needs the Local Clock although a Local PTP Clock can be present.

A PTP Management Node is a device in the PTP Network that uses PTP management messages to configure and/or monitor PTP Instances within PTP Nodes. A PTP Management Node contains the functionality specified in Clause 15. A PTP Management Node transmits PTP management messages with actions GET, SET, and/or COMMAND, and it receives PTP management messages with actions RESPONSE and/or ACKNOWLEDGE.

NOTE—Manufacturers of devices supporting the PTP protocol typically advertise these devices as “Boundary Clocks”, “Ordinary Clocks”, “Transparent Clocks”, or “Grandmaster Clocks”. These and other product designations are terms outside of this standard. Thus, for example, the term “Boundary Clock” in this standard has a specific meaning defined in the standard and does not necessarily describe a specific “Boundary Clock” product, and absent conformance certification, there is no assurance that a product named a “Boundary Clock” implements the specifications of this standard.

6.2 Principal assumptions about the network and implementation recommendations

The following principal assumptions and recommendations should be followed to ensure correct operation of the protocol. These are discussed in greater detail in later clauses:

- a) PTP assumes that the network eliminates cyclic forwarding of PTP messages within each communication path (e.g., by using a spanning tree protocol). However, the PTP protocol also eliminates cyclic forwarding of PTP messages.
- b) PTP is tolerant of an occasional missed PTP message, duplicated PTP message, or PTP message that arrived out of order. However, PTP assumes that such impairments are rare.
- c) PTP was designed assuming a multicast communication model. PTP also supports a unicast communication model, and a mixed multicast/unicast communication model, if the behavior of the protocol is preserved. PTP assumes that Announce messages are periodically sent by one PTP Port and delivered to all other PTP Ports within a PTP Communication Path. If the PTP Communication Path consists of more than two PTP Ports, the assumption is that Announce messages are either sent in multicast or the Announce information is replicated to all PTP Ports in the PTP Communication Path using unicast PTP messages. PTP Ports discover other ports within a PTP Communication Path through the receipt of multicast Announce messages. When multicast communication is not available, another form of discovery (e.g., by configuration) is required; see, for example, 17.4. PTP management messages sent to all PTP Ports also require either multicast messaging or replication of the PTP management message to all PTP Ports within the PTP Communication Path.
- d) If two-step PTP Ports are used, then the PTP Network must be designed such that the general message takes the same path (Sync and Follow_Up), or in some cases the reverse path (Delay_Req and Delay_Resp), as the event message through a Transparent Clock. Failure to do this will result in a condition where PTP does not calculate path delay properly. This condition is undetectable and may introduce additional jitter and wander, but it will not cause the operation of the PTP protocol to cease.
- e) PTP assumes that the number of Boundary Clocks forming the master–slave synchronization hierarchy from the Grandmaster PTP Instance to any Slave PTP Instance is less than 255 (see 9.3.2.5).
- f) Network components that are not PTP-aware, for example, bridges, introduce timing jitter and wander that, if uncorrected, can degrade time transfer accuracy. Since the jitter and wander are often traffic dependent, the network traffic patterns should be designed to minimize the traffic and to minimize the variation in the traffic load. It is also recommended that PTP event messages be sent in high priority compared with other data (see A.5.3.3). Whenever possible such network components should be replaced by PTP Boundary Clocks or Transparent Clocks.
- g) The underlying transport protocol is structured such that a message timestamp point can be defined.

6.3 PTP Networks

The PTP Instances in a PTP Network execute the PTP protocol to synchronize the clocks in the PTP Instances.

The protocol executes within a logical scope called a “domain” as specified in 7.1. Unless otherwise specified, all PTP messages, data sets, state machines, and all other PTP entities are always associated with a particular domain. A PTP Network can support multiple domains. Within this standard, the time established within one domain by the protocol is independent of the time in other domains.

The PTP protocol is a distributed protocol that specifies how the real-time clocks in the PTP Instances of each domain of the PTP Network synchronize with each other. The PTP Instances containing these clocks

are organized into a master–slave synchronization hierarchy whereby the PTP Instance at the top of the hierarchy—the Grandmaster PTP Instance—determines the reference time for the entire domain. The synchronization is achieved by exchanging PTP timing messages with the Slave PTP Instances using the timing information to adjust their clocks to the time of their respective Parent PTP Instances in the hierarchy.

PTP Instances in a PTP Network communicate with each other via a communication network. The PTP Network can include translation devices between segments implementing different network communication protocols.

6.4 PTP message classes

The protocol defines event and general PTP messages. PTP event messages are timed messages in that an accurate timestamp is generated at both transmission and receipt as described in 6.6.5. PTP general messages do not require accurate timestamps.

The set of PTP event messages consists of:

- a) Sync (see 13.6)
- b) Delay_Req (see 13.6)
- c) Pdelay_Req (see 13.9)
- d) Pdelay_Resp (see 13.10)

The set of PTP general messages consists of:

- Announce (see 13.5)
- Follow_Up (see 13.7)
- Delay_Resp (see 13.8)
- Pdelay_Resp_Follow_Up (see 13.11)
- Management (see Clause 15)
- Signaling (see 13.12)

PTP timing messages are the set of all PTP event messages plus the following PTP general messages: Follow_Up, Delay_Resp, and Pdelay_Resp_Follow_Up.

The Sync, Delay_Req, Follow_Up, and Delay_Resp messages are used to generate and communicate the timing information needed to synchronize Ordinary Clocks and Boundary Clocks using the delay request–response mechanism.

The Pdelay_Req, Pdelay_Resp, and Pdelay_Resp_Follow_Up messages are used to measure the PTP Link delay between two PTP Ports implementing the peer-to-peer delay mechanism. The PTP Link delay is used to correct timing information in Sync and Follow_Up messages transmitted over PTP Links. Ordinary Clocks and Boundary Clocks that implement the peer-to-peer delay mechanism can synchronize using the measured PTP Link delays and the information in the Sync and Follow_Up messages.

The Announce message is used to establish the synchronization hierarchy.

The PTP management messages are used to query and update the PTP data sets maintained by PTP Instances. These messages are also used to configure a PTP Network and for initialization and fault management. PTP management messages are used between PTP Management Nodes and PTP Instances.

NOTE—PTP Nodes can support other management mechanisms as an alternative to PTP management messages (see 8.1.4.1).

The Signaling messages are used for communication between PTP Instances for all other purposes. For example, Signaling messages can be used for negotiation of the rate of unicast PTP messages between a Master PTP Instance and its Slave PTP Instances.

All PTP Messages can be extended by means of a standard type, length, value (TLV) extension mechanism. For example, the PATH_TRACE message extensions can be used to detect rogue Announce messages. See 16.2.1 for more detail on rogue Announce messages.

6.5 PTP device types

6.5.1 General

There are five basic types of PTP devices, as follows:

- a) Ordinary Clock
- b) Boundary Clock
- c) End-to-end Transparent Clock
- d) Peer-to-peer Transparent Clock
- e) PTP Management Node (6.1.3)

All five types implement one or more aspects of the protocol.

Two mechanisms are used in PTP to measure the propagation delay between PTP Ports. The first one, the delay request-response mechanism, uses the messages Sync, Delay_Req, Delay_Resp, and if required, Follow_Up (see 11.3). The second one, the peer-to-peer delay mechanism, uses the messages Pdelay_Req, Pdelay_Resp, and if required, Pdelay_Resp_Follow_Up (see 11.4). PTP Ports on Ordinary Clocks and Boundary Clocks can be implemented using either mechanism. PTP Ports on end-to-end Transparent Clocks participate in the delay request-response mechanism. PTP Ports on peer-to-peer Transparent Clocks use the peer-to-peer delay mechanism. These two mechanisms do not interwork on the same PTP Communication Path. In addition, the peer-to-peer delay mechanism is restricted to topologies where each peer-to-peer PTP Port communicates PTP messages with at most one other peer-to-peer PTP Port (see 11.4.3).

The use of the various PTP Instance types is limited as follows:

- Ordinary Clocks and Boundary Clocks with PTP Ports implementing only the peer-to-peer delay mechanism, and peer-to-peer Transparent Clocks, can only be connected in topologies where each such PTP Port implementing the peer-to-peer delay mechanism communicates PTP messages to and from a single PTP Port also implementing the peer-to-peer delay mechanism (see 11.4.3). Except in carefully designed networks, this will preclude the use of end-to-end Transparent Clocks and bridges that do not support PTP, for example, a conventional bridge, when using the peer-to-peer delay mechanism.
- Ordinary Clocks and Boundary Clocks with PTP Ports implementing only the delay request-response mechanism, and end-to-end Transparent Clocks, can be connected in any topology that excludes other PTP Ports using the peer-to-peer delay mechanism. This precludes the use of peer-to-peer Transparent Clocks in such a PTP Network.
- A Boundary Clock with PTP Ports supporting each of the two mechanisms may be used to bridge between regions supporting the different mechanisms.

6.5.2 Ordinary Clocks and Boundary Clocks

6.5.2.1 Layered model of Ordinary Clocks and Boundary Clocks

6.5.2.1.1 General

Figure 5 illustrates the layered model of an Ordinary Clock, whereas Figure 6 illustrates the layered model of a Boundary Clock.

The purpose of these models is to clarify the relationships and functions of the clauses of this standard. These models are not to be construed as implementation specifications.

NOTE—The models of Figure 5 and Figure 6 are restatements of Figures 2 and 3 of IEEE Std 1588-2008 except for the model of timestamp generation, which has been moved to 6.5.2.1.4. With the exception of the Special Port for media not using PTP timing messages, that is, the rightmost PTP Port in Figure 6, these models are compatible with the models used in IEEE Std 1588-2008.




Figure 5—Layered model of an Ordinary Clock




Figure 6—Layered model of a Boundary Clock

Differences between Ordinary Clocks and Boundary Clocks are elaborated in 6.5.2.2 and 6.5.2.3, respectively. The principal features of the models illustrated in Figure 5 and Figure 6 are as follows.

Each figure is divided into three regions:

- a) The external environment shown at the top of the figures. The main goal of PTP is to transfer time from an application-specific source of time to application devices. As noted in 6.3, the PTP protocol executes within a logical scope called a domain. The external environment is not within the domain as indicated by the solid line annotated as external and internal to the domain. Contained in the external environment are
 - 1) Clock Sources: External sources of the time to be distributed by PTP, for example, GNSS receivers.
 - 2) Clock Sinks: Users of the PTP-distributed time. Typical users are sensors, actuators, and computation elements.
- b) PTP Instance.
- c) Communication network over which a PTP Instance communicates with its peers in the domain.

6.5.2.1.2 External environment and the Source Dependent and Sink Dependent blocks

6.5.2.1.2.1 Introducing time into a PTP Instance

Time from one or more external Clock Sources, that is, Clock Source blocks, can be introduced into a PTP Instance (and the domain) only via Source Dependent blocks. This is a one-way transfer of time from the external environment into the PTP Instance, that is, the operation of PTP does not change the time maintained by the external Clock Source. This transfer can require timing control signals as illustrated.

Typical external Clock Sources are GNSS receivers or terrestrial radio links. See Table 6 for an enumeration used to characterize such sources. Source Adapter blocks manage the differences between an interface unique to the external source and a common implementation-specific interface to the External Source block. The External Source block is responsible for selecting which of the external Clock Sources, or if none are available an internal source, is to be used as the basis for assigning values to the members of the default data set. In the event that the PTP Instance is selected as the Grandmaster PTP Instance, the External Source block also provides data to populate the time properties data set. See 7.6.6 for specifications on updating these data sets.

6.5.2.1.2.2 Providing time to an external application

Time from a PTP Instance can be provided to an external application, that is, Clock Sink block, only via the Sink Dependent block (see 7.6.7). This is a one-way transfer of time from the PTP Instance to external Clock Sinks; that is, the operation of PTP is not changed as a result of this time transfer. This transfer can require timing control signals as illustrated. There can be multiple sinks served by a PTP Instance.

Typical external applications, that is, Clock Sinks, are sensors, actuators, and waveform generators. Sink Adapter blocks manage the differences between an interface unique to the external Clock Sink and a common implementation-specific interface to the External Sink block. The External Sink block is responsible for delivery of either PTP- or ARB-based time (see 7.2.1), as well as for accompanying meta-data to this interface.

6.5.2.1.3 Common core, data sets, internal clocks, and management

6.5.2.1.3.1 PTP Common Core block

The PTP Common Core block is responsible for all operations of the PTP protocol that are common to the other blocks to the PTP Instance as a whole, as well as for coordinating the activities of the other blocks in the PTP Instance. These responsibilities include the following:

- a) Operations, except for management (see Clause 15), of the default, parent, current, time properties, and optional data sets that are not per PTP Port (see Clause 8)
- b) Aspects of state configuration, including the BMCA, that require considerations of information from all PTP Ports (see Clause 9)
- c) Aspects of PTP message handling that require considerations of information from all PTP Ports, and inter-port PTP message and information handling for retransmission (see Clause 10 and Clause 11)
- d) Aspects of optional PTP features that require consideration of information from all PTP Ports (see Clause 16 and Clause 17, Annex J, and Annex P)

6.5.2.1.3.2 Clocks associated with a PTP Instance

The PTP protocol maintains time within each PTP Instance and within the domain, based on either the timescale PTP or the timescale ARB (see 7.2.1). The protocol can distribute information enabling an external Clock Sink to convert from one of these timescales to an application-specific timescale. This conversion can be done in the Sink Adapter block or by the external Clock Sink itself.

The timescale PTP or the timescale ARB is maintained in the Local PTP Clock block. The time of Local PTP Clock represents the local estimate of the time of the Grandmaster Clock, that is, the time of the Local PTP Clock of the Grandmaster PTP Instance of the domain. There are two methods of maintaining the Local PTP Clock:

- 1) Maintain an actual physical clock, that is, a local oscillator and a time counter, synchronized to the time of the Grandmaster Clock, or in the case of the Grandmaster Clock to the selected external time source represented in Figure 5 and Figure 6 by the Clock Source block. In this method, the Local PTP Clock is identical to the Local Clock.
- 2) Maintain a model that computes the timescale PTP or the timescale ARB based on the Local Clock. In this method, the model incorporates the results of the synchronization procedures.

The synchronization procedures are discussed in Clause 11 and Clause 12.

All blocks in the PTP Instance can obtain time from these clocks.

6.5.2.1.3.3 Management

The Management block (labeled “Mgmt” in Figure 5 and Figure 6) can communicate with all blocks in the PTP Instance. The Management block allows observing and, in some cases, configuring aspects of the PTP data sets and the operation of the protocol. Various management mechanisms can be applied to the Management block (see 8.1.4.1).

PTP management messages (see Clause 15) use the targetPortIdentity field of each message to identify the specific PTP Instance (and PTP Port) that is managed by the message. Therefore, for PTP management messages, the Management block operates within each PTP Instance.

Management mechanisms (e.g., SNMP) other than PTP management messages are typically implemented in the PTP Node. Following the organization specified in 8.1.4.2, these management mechanisms provide a data model for access to each PTP Instance (and PTP Port). Therefore, for these management mechanisms, the Management block operates within the PTP Node, but it provides access to each PTP Instance through its data model.

6.5.2.1.4 PTP Ports and the PTP Network

The models of this subclause apply to Boundary Clocks, Ordinary Clocks, and Transparent Clocks, except that Transparent Clocks do not have Special Ports.

Each PTP Port consists of a PTP Port block and a Media Dependent block.

The PTP Port block is separated from the Media Dependent block by the MDMI interface. The MDMI interface must be supported for Special Ports and can be supported for E2E and P2P PTP Ports. The MDMI interface for Special Ports is specified in 11.5.

NOTE 1—The MDMI interface is new to this edition and is introduced to facilitate the use within the PTP protocol of media that implement their own methodology for measuring path delay and transferring time, for example, IEEE Std 802.11.

The PTP Port block is the media-independent portion of the PTP Port that is responsible for all per-port functionality that does not depend on the communication media or the time transfer method, for example, E2E, P2P or Special, being used. These responsibilities include:

- Operations and maintenance of the PTP Port data sets and the <foreignMasterList> (see 8.2.15 and 9.3.2.4, respectively), and any optional data sets, if present, that are per PTP Port
- Operation of the per-port portion of the BMCA (see 9.3)

The Media Dependent block is the media-dependent portion of the PTP Port that consists of two sub-blocks: a PTP MD Adapter and a Network Interface Stack. The PTP MD Adapter block will differ depending on the time transfer methodology used, that is, E2E, P2P, or Special.

The MD Adapters and Network Interface Stack blocks are separated by the Network Interface. The PTP MD Adapter is responsible for media-dependent functions of the PTP protocol, including:

- Handling the transmission and reception of all PTP messages
- Handling the transfer of time between the Network Interface Stack and the PTP Port block

The PTP MD Adapter block communicates with the Network Interface Stack via three interfaces: General Message Functions, Event Message Functions, and Other Adapter Functions, as illustrated in Figure 7, all on a single physical port.

The Event Message Function block is used to send and receive PTP event messages, which are timestamped by the Timestamp Generation block based on the value of the Timestamping Clock. To achieve the best accuracy and least jitter, the detection of PTP event messages and subsequent capture of the time of the Timestamping Clock is preferably implemented as close to the network as possible. Many of the modern PHY chips therefore generate a signal upon detecting a PTP event messages. This signal is then used to capture the timestamp. Therefore, in Figure 7, the detection of PTP event messages and the generation of the timestamp can occur either in the MD Adapter block or, preferably, in the Network Interface Stack as illustrated. See 6.6.5 for a more complete discussion of timestamp generation. In either case, the MD Adapter block is responsible for transferring both the information contained in the PTP event message as well as the timestamp to the PTP Port block.

The General Message Function block is used to send and receive PTP general messages. The Other Adapter Functions block handles all other functions and any coordination of the other blocks of the MD Adapter block.

There are three forms of the PTP Port:

- a) The end-to-end (E2E) form consists of a PTP Port block and a corresponding media-dependent E2E PTP MD Adapter block. These blocks implement the end-to-end methodology of time transfer and determining path delay (see 11.3). These blocks make use of PTP messages communicated over the PTP Network to other PTP Instances in the domain to implement two-way time transfer between PTP Instances.
- b) The peer-to-peer (P2P) form consists of a PTP Port block and a corresponding media-dependent P2P PTP MD Adapter block. These blocks implement the peer-to-peer methodology of time transfer and determining path delay (see 11.4). These blocks make use of PTP messages communicated over the PTP Network to other PTP Instances in the domain to implement two-way time transfer between PTP Instances.
- c) The Special form consists of a PTP Port block and a corresponding media-dependent Special PTP MD Adapter block. These blocks implement time transfer over media using media-specific time transfer methods (not PTP timing messages) to transfer time between PTP Instances in a domain (see 11.5, which specifies the details of the Special Ports including the MDMI interface).

NOTE 2—The Special Port allows the use of media-specific time transfer methods specified by media standards organizations, for example, IEEE Std 802.11. See 11.5.




Figure 7—Model of the message interfaces of the MD Adapter block

6.5.2.2 Properties of an Ordinary Clock

An Ordinary Clock in a domain supports a single copy of the protocol and has a single PTP Port (see Figure 5). This single PTP Port can be an E2E, a P2P, or a Special Port see (6.5.2.1.4). The Ordinary Clock can be the Grandmaster PTP Instance, or it can be a Slave PTP Instance in the master–slave hierarchy.

In some applications, such as industrial automation, an Ordinary Clock can also be associated with an application device such as a sensor or an actuator. In a telecommunications application, an Ordinary Clock can be associated with a timing demarcation device. Using the Sink Dependent blocks, the Local PTP Clock provides time support to one or more such a device as illustrated in Figure 5.

6.5.2.3 Properties of a Boundary Clock

The Boundary Clock has several PTP Ports (see Figure 6). Each PTP Port of a Boundary Clock is like the PTP Port of an Ordinary Clock.

All PTP messages except PTP management messages are terminated at a Boundary Clock and are not retransmitted. PTP management messages are retransmitted by other PTP Ports on the Boundary Clock subject to restrictions to limit the propagation of these PTP messages within the PTP Network.

The Boundary Clock model of Figure 6 is applicable only to PTP messages and non-PTP messages used by the Special Port.

A Boundary Clock can become the Grandmaster PTP Instance of the domain.

6.5.3 End-to-end Transparent Clocks

The model of an end-to-end Transparent Clock is illustrated in Figure 8.




Figure 8—Layered model of an End-to-end Transparent Clock