

INTERNATIONAL STANDARD



**Industrial communication networks – High availability automation networks –
Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless
Redundancy (HSR)**



THIS PUBLICATION IS COPYRIGHT PROTECTED

Copyright © 2016 IEC, Geneva, Switzerland

All rights reserved. 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 either IEC or IEC's member National Committee in the country of the requester. If you have any questions about IEC copyright or have an enquiry about obtaining additional rights to this publication, please contact the address below or your local IEC member National Committee for further information.

IEC Central Office
3, rue de Varembe
CH-1211 Geneva 20
Switzerland

Tel.: +41 22 919 02 11
Fax: +41 22 919 03 00
info@iec.ch
www.iec.ch

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 corrigenda or an amendment might have been published.

IEC Catalogue - webstore.iec.ch/catalogue

The stand-alone application for consulting the entire bibliographical information on IEC International Standards, Technical Specifications, Technical Reports and other documents. Available for PC, Mac OS, Android Tablets and iPad.

IEC publications search - www.iec.ch/searchpub

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 It also gives information on projects, replaced and withdrawn publications.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 also once a month by email.

Electropedia - www.electropedia.org

The world's leading online dictionary of electronic and electrical terms containing 20 000 terms and definitions in English and French, with equivalent terms in 15 additional languages. Also known as the International Electrotechnical Vocabulary (IEV) online.

IEC Glossary - std.iec.ch/glossary

65 000 electrotechnical terminology entries in English and French extracted from the Terms and Definitions clause of IEC publications issued since 2002. Some entries have been collected from earlier publications of IEC TC 37, 77, 86 and CISPR.

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: csc@iec.ch.

IECNORM.COM : Click to view the full text of IEC 60332-3:2016 RVV

INTERNATIONAL STANDARD



**Industrial communication networks – High availability automation networks –
Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless
Redundancy (HSR)**

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

ICS 25.040, 35.040

ISBN 978-2-8322-3291-0

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

CONTENTS

FOREWORD	8
INTRODUCTION	10
0.1 General	10
0.2 Changes with respect to the previous edition	10
0.3 Patent declaration	11
1 Scope	13
2 Normative references	13
3 Terms, definitions, abbreviations, acronyms, and conventions	14
3.1 Terms and definitions	14
3.2 Abbreviations and acronyms	15
3.3 Conventions	16
4 Parallel Redundancy Protocol (PRP)	16
4.1 PRP principle of operation	16
4.1.1 PRP network topology	16
4.1.2 PRP LANs with linear or bus topology	17
4.1.3 PRP LANs with ring topology	17
4.1.4 DANP node structure	18
4.1.5 PRP attachment of singly attached nodes	19
4.1.6 Compatibility between singly and doubly attached nodes	19
4.1.7 Network management	19
4.1.8 Implication on ~~configuration~~ application	20
4.1.9 Transition to non-redundant networks	20
4.1.10 Duplicate handling	20
4.1.11 Network supervision	26
4.1.12 Redundancy management interface	26
4.2 PRP protocol specifications	26
4.2.1 Installation, configuration and repair guidelines	26
4.2.2 Unicast MAC addresses	27
4.2.3 Multicast MAC addresses	27
4.2.4 IP addresses	27
4.2.5 Nodes	27
4.2.6 Duplicate Accept mode (testing only)	28
4.2.7 Duplicate Discard mode	28
4.3.2 PRP_Supervision frame contents	33
4.3.3 PRP_Supervision frame for RedBox	34
4.3.4 Reception of a PRP_Supervision frame and NodesTable	34
4.3 PRP_Supervision frame	32
4.3.1 PRP_Supervision frame format	32
4.4 Bridging node	35
4.5 Constants	35
4.6 PRP service specification	35
5 High-availability Seamless Redundancy (HSR)	36
5.1 HSR objectives	36
5.2 HSR principle of operation	36
5.2.1 Basic operation with a ring topology	36

5.2.2	DANH node structure	38
5.2.3	Topology	38
5.2.4	RedBox structure	46
5.3	HSR node specifications	48
5.3.1	HSR operation	48
5.3.2	DANH receiving from its link layer interface	48
5.3.3	DANH receiving from an HSR port	49
5.3.4	DANH forwarding rules	50
5.3.5	CoS	52
5.3.6	Clock synchronization	52
5.3.7	Deterministic medium access	52
5.4	HSR RedBox specifications	52
5.4.1	RedBox properties	52
5.4.2	RedBox receiving from interlink	53
5.4.3	RedBox forwarding on the ring	55
5.4.4	RedBox receiving from an HSR port	55
5.4.5	RedBox receiving from its link layer interface	57
5.4.6	Redbox ProxyNodeTable handling	57
5.4.7	RedBox CoS	57
5.4.8	RedBox clock synchronization	58
5.4.9	RedBox medium access	58
5.5	QuadBox specification	58
5.6	Duplicate Discard method	58
5.7	Frame format for HSR	58
5.7.1	Frame format for all frames	58
5.7.2	HSR_Supervision frame	59
5.8	Constants	62
5.9	HSR service specification	63
6	Protocol Implementation Conformance Statement (PICS)	64
7	PRP/HSR Management Information Base (MIB)	65
~~Annex A (normative) Use of IEC 61588 and IEEE C37.238 for IEC 62439-3~~		
Annex A (normative) Clocks synchronization over redundant paths in		
IEC 62439-3		96
A.1	Overview	96
A.2	Attachment to redundant LANs by a boundary clock	96
A.3	Attachment to redundant LANs by doubly attached ordinary clocks	97
A.4	PRP mapping to PTP	99
A.4.1	Scenarios and device roles	99
A.4.2	Operation in PRP	101
A.4.3	Configuration specification	102
A.4.4	Specifications of DANP as DAC	103
A.4.5	Clock model of a RedBox for PTP	103
A.5	HSR Mapping to PTP	120
A.5.1	PTP traffic in HSR	120
A.5.2	HSR nodes specifications	123
A.5.3	Redundant clocks in HSR	124
A.5.4	Attachment of an MC to an external LAN	124
A.6	PRP to HSR Mapping	125
A.6.1	Connection methods	125

A.6.2	PRP-HSR connection by BC	125
A.6.3	PRP-HSR connection by TCs	126
A.7	Doubly attached clock model	127
A.7.1	State machine	127
A.7.2	Supervision of the port	130
A.7.3	BMCA for paired ports	131
A.7.4	Selection of the port state	132
A.8	PTP datasets for high availability	132
A.8.1	General	132
A.8.2	Data types	132
A.8.3	Datasets for ordinary or boundary clocks	133
A.8.4	Object for transparent clocks	137
~~Annex B (informative) Deterministic medium access in HSR~~		
Annex B (normative) PTP profile for Power Utility Automation – Redundant clock attachment		141
B.1	Application domain	141
B.2	PTP profile specification	141
B.3	Redundant clock attachment	141
Annex C (normative) PTP profiles for high-availability automation networks		142
C.1	Application domain	142
C.2	PTP profile specification	142
C.3	Clock types	142
C.4	Protocol specification common	143
C.5	Protocol specification for L3E2E automation profile	143
C.6	Protocol specification for L2P2P automation profile	143
C.7	Timing requirements	144
C.7.1	Measurement conditions	144
C.7.2	Network time inaccuracy	144
C.7.3	Network elements	144
C.7.4	Requirements for grandmasters	144
C.7.5	Requirements for TCs	145
C.7.6	Requirements for BCs	145
C.7.7	Requirements for media converters	145
C.7.8	Requirements for links	145
C.8	Network engineering	146
C.9	Default settings	146
C.10	Redundant clock handling	147
C.11	Protocol Implementation Conformance Statement (PICS)	148
C.11.1	Conventions	148
C.11.2	PICS	148
Annex D (informative) Precision Time Protocol tutorial for IEC 62439-3		150
D.1	Objective	150
D.2	Precision and accuracy	150
D.3	PTP clock types	151
D.4	PTP main options	152
D.5	Layer 2 and layer 3 communication	153
D.6	1-step and 2-step correction	153
D.6.1	Time correction in TCs	153
D.6.2	2-step to 1-step translation	154

D.7	End-To-End link delay measurement	156
D.7.1	General method	156
D.7.2	End-to-End link delay measurement with 1-step clock correction	156
D.7.3	End-to-End link delay measurement with 2-step clock correction	157
D.7.4	End-to-End link delay calculation by Delay_Req/Delay_Resp	158
D.8	Peer-to-Peer link delay calculation	158
D.8.1	Peer-to-Peer link delay calculation with 1-step correction	158
D.8.2	Peer-to-Peer link delay calculation with 2-step correction	159
Annex E (normative)	Management Information base for singly and doubly attached clocks.....	161
Bibliography		186
Figure 1	– PRP example of general redundant network.....	16
Figure 2	– PRP example of redundant network as two LANs (bus topology).....	17
Figure 3	– PRP example of redundant ring with SANs and DANPs	18
Figure 4	– PRP with two DANPs communicating.....	18
Figure 5	– PRP RedBox, transition from single to double LAN.....	20
Figure 6	– PRP frame extended by an RCT	21
Figure 7	– PRP VLAN-tagged frame extended by an RCT	22
Figure 8	– PRP padded frame closed by an RCT	22
Figure 9	– Duplicate Discard algorithm boundaries	24
Figure 10	– HSR example of ring configuration for multicast traffic.....	36
Figure 11	– HSR example of ring configuration for unicast traffic	37
Figure 12	– HSR structure of a DANH	38
Figure 13	– HSR example of topology using two independent networks.....	39
Figure 14	– HSR example of peer coupling of two rings	40
Figure 15	– HSR example of connected rings	41
Figure 16	– HSR example of coupling two redundant PRP LANs to a ring	42
Figure 17	– HSR example of coupling from a ring node to redundant PRP LANs	43
Figure 18	– HSR example of coupling from a ring to two PRP LANs	44
Figure 19	– HSR example of coupling three rings to one PRP LAN.....	45
Figure 20	– HSR example of meshed topology	46
Figure 21	– HSR structure of a RedBox.....	47
Figure 22	– HSR frame without a VLAN tag	58
Figure 23	– HSR frame with VLAN tag.....	59
Figure 24	– HSR node with management counters.....	63
Figure 25	– HSR RedBox with management counters	64
Figure A.1	– Doubly Attached Clock as BC (MCA is best master).....	96
Figure A.2	– Doubly Attached Clock when MCA is best master	98
Figure A.3	– Doubly attached clocks when OC1 is best master	99
Figure A.4	– Elements of PRP networks	101
Figure A.5	– Connection of a master clock to an ordinary clock over PRP	102
Figure A.6	– PRP RedBox as BCs (OC3 and BC7 are best masters)	104
Figure A.7	– RedBox DABC clock model	105

Figure A.8 – PRP RedBoxes as DABC with E2E – BC7 is master	106
Figure A.9 – PRP RedBoxes as DABC with E2E – timing	107
Figure A.10 – PRP RedBoxes as DABC with P2P – OC5 is best master	108
Figure A.11 – PRP RedBoxes as DABC with P2P – timing	109
Figure A.12 – PRP RedBox as DATC with E2E –signal flow	110
Figure A.13 – PRP RedBox as DATC with E2E – timing	112
Figure A.14 – PRP RedBox as DATC with P2P	113
Figure A.15 – PRP RedBox as DATC with P2P – timing	114
Figure A.16 – PRP RedBox as SLTC with E2E	117
Figure A.17 – PRP RedBox as SLTC with E2E – timing	118
Figure A.18 – PRP RedBox as SLTC with P2P	119
Figure A.19 – HSR with one GMC	121
Figure A.20 – PTP messages sent and received by an HSR node (1-step)	122
Figure A.21 – PTP messages sent and received by an HSR node (2-step)	123
Figure A.22 – Attachment of a GMC to an HSR ring through a RedBox as TC	125
Figure A.23 – PRP to HSR coupling by BCs	126
Figure A.24 – PRP to HSR coupling by TCs	127
Figure A.25 – Port states including transitions for redundant operation	128
Figure A.26 – BMCA for redundant masters	131
Figure D.1 – Precision and accuracy example	150
Figure D.2 – Precision Time Protocol principle	151
Figure D.3 – Precision Time Protocol elements	152
Figure D.4 – Delays and time-stamping logic in TCs	153
Figure D.5 – Correction of the Sync message by 1-step and 2-step (peer-to-peer)	154
Figure D.6 – Translation from 2-step to 1-step in TCs	155
Figure D.7 – Translation from 2-step to 1-step – message view	156
Figure D.8 – End-to-end link delay measurement with 1-step clock correction	157
Figure D.9 – End-to-end delay measurement with 2-step clock correction	158
Figure D.10 – Peer-to-peer link delay measurement with 1-step clock correction	159
Figure D.11 – Peer-to-peer link delay measurement with 2-step clock correction	160
Table 1 – Duplicate discard cases	25
Table 2 – Monitoring data set	28
Table 3 – NodesTable attributes	29
Table 4 – PRP_Supervision frame with no VLAN tag	32
Table 5 – PRP_Supervision frame with (optional) VLAN tag	33
Table 6 – PRP_Supervision frame contents	34
Table 7 – PRP_Supervision TLV for Redbox	34
Table 8 – PRP constants	35
Table 9 – HSR_Supervision frame with no VLAN tag	60
Table 10 – HSR_Supervision frame with optional VLAN tag	61
Table 11 – HSR Constants	63
Table A.1 – States	129

Table A.2 – Transitions	130
Table A.3 – Variables	130
Table C.1 – PTP attributes for the Industrial Automation profile	147
Table C.2 – PICS for clocks	149

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

INTERNATIONAL ELECTROTECHNICAL COMMISSION

**INDUSTRIAL COMMUNICATION NETWORKS –
HIGH AVAILABILITY AUTOMATION NETWORKS –****Part 3: Parallel Redundancy Protocol (PRP) and
High-availability Seamless Redundancy (HSR)**

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 Publication(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. IEC collaborates closely with the International Organization for Standardization (ISO) in accordance with conditions determined by agreement between the two organizations.
- 2) The formal decisions or agreements 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.
- 3) IEC Publications have the form of recommendations for international use and are accepted by IEC National Committees in that sense. While all reasonable efforts are made to ensure that the technical content of IEC Publications is accurate, IEC 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 transparently to the maximum extent possible in their national and regional publications. Any divergence between any IEC Publication and the corresponding national or regional publication shall be clearly indicated in the latter.
- 5) IEC itself does not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. IEC is 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 its directors, employees, servants or agents including individual experts and members of its technical committees and IEC National Committees 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 Publication or any other IEC 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.

This redline version of the official IEC Standard allows the user to identify the changes made to the previous edition. A vertical bar appears in the margin wherever a change has been made. Additions are in green text, deletions are in strikethrough red text.

International Standard IEC 62439-3 has been prepared by subcommittee 65C: Industrial networks, of IEC technical committee 65: Industrial-process measurement, control and automation.

This third edition cancels and replaces the second edition published in 2012. This edition constitutes a technical revision.

This edition includes the following significant technical changes with respect to the previous edition:

- a) technical corrections and extension of specifications;
- b) consideration of IEC 61588 clock synchronization with end-to-end delay measurement alongside the existing peer-to-peer delay measurement in PRP.

The text of this standard is based on the following documents:

FDIS	Report on voting
65C/834/FDIS	65C/841/RVD

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

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2.

This International Standard is to be read in conjunction with IEC 62439-1.

A list of all parts in the IEC 62439 series, published under the general title *Industrial communication networks – High availability automation networks*, can be found on the IEC website.

The committee has decided that the contents of this publication will remain unchanged until the stability date indicated on the IEC web site under "<http://webstore.iec.ch>" in the data related to the specific publication. At this date, the publication will be

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

IMPORTANT – The “colour inside” logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this publication using a colour printer.

INTRODUCTION

0.1 General

IEC 62439-3 belongs to the IEC 62439 series “*Industrial communication networks – High availability automation networks*”. It specifies the PRP and HSR **seamless** redundancy protocols. It was adopted by IEC TC57 WG10 as the redundancy method for demanding substation automation networks ~~based on IEC 61850 series, introducing new requirements operating on layer 2 networks, according to IEC 61850-8-1 and IEC 61850-9-2.~~

The **seamless** redundancy principle has been extended to clocks operating according to the Precision Time Protocol (IEC 61588) and attached to redundant networks. Two variants are specified: L3E2E for clocks which operate on layer 3 networks with end-to-end link delay measurement (E2E) and L2P2P for clocks that operate on layer 2 with peer-to-peer link delay measurement (P2P).

0.2 Changes with respect to the previous edition

~~The major changes with respect to IEC 62439-3:2010 are listed below.~~

~~Aligning the sequence number between PRP and HSR, to enable coupling of HSR and PRP networks and simplify the implementation of dual mode nodes in hardware. At the same time, introduce a suffix in the PRP Redundancy Control Trailer to allow better identification, future extensions and coexistence with other protocols that also happen to use a trailer. This change is not backwards compatible, so means are provided to identify the version and ensure that the networks are homogeneous.~~

~~Removing all implementation restrictions on the Duplicate Discard algorithm (especially references to the drop window algorithm and references to connection orientation) since other methods such as hash tables can be used.~~

~~Removing the purging of the duplicate table. Replace this specific method by requiring that any Duplicate Discard algorithm provides a mechanism to remove old entries, thus ensuring that a node can properly reboot.~~

~~Making node tables optional for simple nodes to simplify hardware implementation.~~

~~Suppression of explicit mention of the HSR-PRP mode (PRP with HSR Tags), but allow it through the Mode N (no forwarding).~~

~~Introducing Mode T (forward through) to allow maintenance laptops to configure an open ring when attached to one end and Mode M (mixed) to allow forwarding of non HSR tagged frames in a closed ring.~~

~~Recommending the position of connectors, rather than impose it.~~

~~Defining the behaviour of an HSR node when non HSR frames are encountered without requiring the recording of the source addresses and specify how IEEE 802.1D:2004, Table 7-10 frames are treated.~~

~~Prefixing the supervision frames on HSR by an HSR tag to simplify the hardware implementation and introduce a unique EtherType for HSR to simplify processing.~~

~~Changing the rule for the RedBox to allow more than one PRP network to be connected to an HSR ring, and introduce an identifier per RedBox pair.~~

~~Specifying tagging of IEC 61588 frames to follow IEEE C37.238 recommendations (informal).~~

~~Suppressing MAC address substitution.~~

~~Adapting the MIB to above changes.~~

The major changes with respect to IEC 62439-3:2012 are:

- Subclause 4.1.10.3 has been rewritten to explain the calculation of the duplicate rejection for different speeds.
- Annex A has been redrafted as a general concept for doubly attached clocks applicable to end-to-end (E2E) and to peer-to-peer (P2P) link delay measurement; the principle of paired port operation has now been specified in terms of a state machine based on IEC 61588:2009.
- Annex B of IEC 62439-3:2012 has been deleted; its properties are mentioned in 5.3.7.
- Annex B (new) makes the support of redundancy mandatory for IEC/IEEE 61850-9-3 that specifies doubly attached clocks on layer 2, with peer-to-peer delay measurement.
- Annex C specifies two profiles of a precision clock for industrial automation: L3E2E for layer 3, end-to-end delay measurement and L2P2P for layer 2, peer-to-peer delay measurement.
- Annex D contains the tutorial information on IEC 61588:2009 for understanding the above annexes. It was contained in IEC 62439-3:2012 Annex A.
- Annex E (MIB) contains the SNMP Management Information Base to be used for singly and doubly attached clocks in all profiles.

0.3 Patent declaration

The International Electrotechnical Commission (IEC) draws attention to the fact that it is claimed that compliance with this document may involve the use of patents concerning filtering of redundant frames in a network node (Siemens Aktiengesellschaft – EP 2127329, US 8184650, CN 101611615B) given in 5.2.3.3.

IEC takes no position concerning the evidence, validity and scope of this patent right.

The holder of this patent right has assured the IEC that he/she is willing to negotiate licences under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of this patent right is registered with IEC. Information may be obtained from:

Siemens Aktiengesellschaft
Oto-Hahn-Ring 6
81379 Munich, Germany

The International Electrotechnical Commission (IEC) draws attention to the fact that it is claimed that compliance with this document may involve the use of a patent concerning Reception of redundant and non-redundant frames (ABB ~~Schweiz AG~~ ~~WO 2006/053459 A1~~ Research Ltd – EP 1825657, US ~~20070223533~~ 8582426, CN 101057483, IN 254425) given in 4.2.7, concerning Identifying improper cabling of devices (ABB ~~Schweiz~~ Technology AG – EP ~~2 015 501 A1~~ 2163024, US 8344736, CN 101689985) given in 4.3, concerning Critical device with increased availability (ABB ~~Schweiz AG~~ Research Ltd – EP 2090950 ~~A1~~) given in 4.4, concerning Ring coupling nodes for high availability networks (ABB ~~Schweiz AG~~ Research Ltd ~~WO 2010/010120 A1~~ US 8582424, EP 2327185, CN 102106121) given in 5.2.3.

IEC takes no position concerning the evidence, validity and scope of these patent rights.

The holder of these patent rights has assured the IEC that he/she is willing to negotiate licences under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of these patent rights is registered with IEC. Information may be obtained from:

ABB Schweiz AG
Intellectual Property CH-IP (CH-150016-L)
Brown Boveri Strasse 6
CH-5400 Baden, Switzerland
ch-ip.patent@abb.com

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights other than those identified above. IEC shall not be held responsible for identifying any or all such patent rights.

ISO (www.iso.org/patents) and IEC (<http://patents.iec.ch>) maintain on-line data bases of patents relevant to their standards. Users are encouraged to consult the data bases for the most up to date information concerning patents.

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

INDUSTRIAL COMMUNICATION NETWORKS – HIGH AVAILABILITY AUTOMATION NETWORKS –

Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

1 Scope

The IEC 62439 series is applicable to high-availability automation networks based on the ~~ISO/IEC 8802-3~~ Ethernet technology.

This part of IEC 62439 specifies two redundancy protocols designed to provide seamless recovery in case of single failure of an inter-bridge link or bridge in the network, which are based on the same scheme: ~~duplication of the LAN, resp. duplication of the transmitted~~ parallel transmission of duplicated information.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

IEC 60050-191, *International Electrotechnical Vocabulary – Chapter 191: Dependability and quality of service*

IEC 61588:2009, *Precision clock synchronization protocol for networked measurement and control systems*

IEC TR 61850-90-4:2013, *Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines*

IEC 62439-1, *Industrial communication networks – High availability automation networks – Part 1: General concepts and calculation methods*

~~IEC 62439-2, *Industrial Communication networks – High availability automation networks – Part 2: Media Redundancy Protocol (MRP)*~~

~~IEC 62439-6, *Industrial communication networks – High availability automation networks – Part 6: Distributed Redundancy Protocol (DRP)*~~

~~IEC 62439-7, *Industrial communication networks – High availability automation networks – Part 7: Ring-based Redundancy Protocol (RRP)*~~

IEC/IEEE 61850-9-3:—, *Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation (proposed IEC 61850-9-3)*¹

¹ To be published.

~~ISO/IEC 8802-3:2000, Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications~~

ISO/IEC/IEEE 8802-3:2014, *Standard for Ethernet*

IEEE 802.1D:2004, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges*

IEEE 802.1Q:2011 2014, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges and Virtual Bridge Local Area Network*

IETF RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

IETF RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

3 Terms, definitions, abbreviations, acronyms, and conventions

3.1 Terms and definitions

For the purposes of this document, the terms and definitions given in IEC 60050-191 and in IEC 62439-1, as well as the following, apply.

3.1.1

grandmaster-capable

ordinary clock or boundary clock that is able to take the role of a grandmaster

Note 1 to entry: A grandmaster-capable clock is not necessarily connected to a recognized time source.

3.1.2

interlink

link that connects two network hierarchies

3.1.3

network time inaccuracy

time inaccuracy evaluated or measured between a particular slave clock and the time reference signal of its grandmaster clock(s), considering the worst network topology

3.1.4

RedBox

device attaching singly attached nodes to a redundant network

3.1.5

QuadBox

quadruple port device connecting two peer HSR rings, which behaves as an HSR node in each ring and is able to filter the traffic and forward it from ring to ring

3.1.6

HSR frame

frame that carries as EtherType the HSR_ethertype

3.1.7

switching logic

hardware that transmits a frame from one port to another port, possibly providing cut-through

3.1.8**time error**

deviation from the time reference used for measurement or synchronization, evaluated over a short time span

3.1.9**time inaccuracy**

time error not exceeded by 99,7% of the measurements, evaluated over a series of 1 000 measurements (about 20 min) in steady state

Note 1 to entry: In this document, the words “bridge” and “bridging” are synonymous with the words “switch” and “switching” respectively, when they apply to layer 2 connectivity.

3.2 Abbreviations and acronyms

For the purposes of this document, the abbreviations and acronyms given in IEC 62439-1, as well as the following apply.

ARP	Address Resolution Protocol
BC	Boundary Clock, as defined in IEC 61588:2009
BMCA	Best Master Clock Algorithm, as defined in IEC 61588:2009
BOC	Boundary Ordinary Clock
CoS	Class of Service (layer 2 quality of service)
DAC	Doubly Attached Clock
DABC	Doubly Attached Boundary Clock
DAN	Doubly Attached Node
DANH	Doubly Attached Node implementing HSR
DANP	Doubly Attached Node implementing PRP
DATC	Doubly Attached Transparent Clock
E2E	End-to-End (delay measurement)
GMC	Grandmaster Clock
HC	Hybrid Clock
HSR	High-availability Seamless Redundancy
IP	Internet Protocol
L2P2P	Layer 2 Peer-to-Peer
L3E2E	Layer 3 End-to-End
LAN	Local Area Network
LLC	Link Layer Control
LRE	Link Redundancy Entity
MAC	Medium Access Control
MC	Master Clock
MIB	Management Information Base
OC	Ordinary Clock, as defined in IEC 61588:2009
P2P	Peer to Peer (delay measurement)
PICS	Protocol Implementation Conformance Statement
PRP	Parallel Redundancy Protocol
PT	Protocol Type
RCT	Redundancy Check Trailer
RSTP	Rapid Spanning Tree Protocol

SAN	Singly Attached Node
SAC	Singly Attached Clock
SLTC	Stateless Transparent Clock
SNMP	Simple Network Management Protocol
TC	Transparent Clock, as defined in IEC 61588:2009
TCP	Transport Control Protocol
UDP	User Data Protocol
VDAN	Virtual Doubly Attached Node (SAN as visible through a RedBox)

3.3 Conventions

This part of the IEC 62439 series follows the conventions defined in IEC 62439-1

4 Parallel Redundancy Protocol (PRP)


4.1 PRP principle of operation

4.1.1 PRP network topology

This redundancy protocol implements redundancy in the nodes rather than in the network (see the definitions in IEC 62439-1), using Doubly Attached Nodes obeying to PRP (DANPs).

A DANP is attached to two independent Local Area Networks (LANs) of similar topology, named LAN_A and LAN_B, which operate in parallel. A source DANP sends the same frame over both LANs and a destination DANP receives it from both LANs within a certain time, consumes the first frame and discards the duplicate.

Figure 1 shows a redundant network consisting of two LANs, each of which can have any topology, e.g. tree, ring or meshed.



IEC

Figure 1 – PRP example of general redundant network


The two LANs are identical in protocol at the MAC-LLC level, but they can differ in performance and topology. Transmission delays may also be different, especially if one of the networks reconfigures itself, e.g. using RSTP, to overcome an internal failure.

The two LANs follow configuration rules that allow the network management protocols such as Address Resolution Protocol (ARP) and Simple Network Management Protocol (SNMP) to operate correctly.

The two LANs have no connection between them and are assumed to be fail-independent. Redundancy can be defeated by single points of failure, such as a common power supply or a direct connection whose failure brings both networks down. Installation guidelines in this document provide guidance to the installer to achieve fail-independence.

4.1.2 PRP LANs with linear or bus topology

As an example of a simpler topology, Figure 2 draws a PRP network as two LANs in linear topology, which may also be a bus topology.



IEC

Figure 2 – PRP example of redundant network as two LANs (bus topology)

4.1.3 PRP LANs with ring topology

The two LANs can have a ring topology, as Figure 3 shows.

NOTE In this case ring Redundancy within each ring is provided by a different mechanism, for instance as defined in IEEE 802.1D or in other IEC 62439 protocols, i.e. MRP.




Figure 3 – PRP example of redundant ring with SANs and DANPs

4.1.4 DANP node structure

Each node has two ports that operate in parallel and that are attached to the same upper layers of the communication stack through the Link Redundancy Entity (LRE), as Figure 4 shows.




Figure 4 – PRP with two DANPs communicating

For the basic communication, the LRE presents toward its upper layers the same interface as a non-redundant network adapter, so the upper layers are unaware of redundancy.

The LRE has two tasks: handling of duplicates and management of redundancy.

When receiving a frame from the node's upper layers, the LRE appends to the frame a Redundancy Check Trailer (RCT) containing a sequence number and sends the frame through both its ports at nearly the same time. The two frames are nearly identical except for the LAN identifier (and the checksum).

The two frames transit through the two LANs with different delays, ideally they arrive at nearly the same time at the destination node.

When receiving frames from the network, the LRE forwards the first received frame of a pair to its node's upper layers and discards the duplicate frame (if it arrives). It removes the RCT if required.

For management of redundancy and checking of the presence of other DANPs, an LRE periodically sends PRP_Supervision frames and can evaluate the PRP_Supervision frames sent by other DANPs.

4.1.5 PRP attachment of singly attached nodes

Singly attached nodes (SANs) can be attached in two ways:

SANs can be attached directly to one LAN only. Such SANs can only communicate with other SANs on the same LAN. For instance, in Figure 1, SAN A1 can communicate with SAN A2, but not with SAN B1 or SAN B2. SANs can communicate (not redundantly) with all DANPs.

SANs can be attached over a Redundancy Box (RedBox) to both LANs, as Figure 1 shows for SAN R1 and SAN R2 (see also 4.1.9). Such SANs can communicate with all DANP and SANs, for instance SAN A1 and SAN R1 can communicate.

NOTE SANs are not aware of PRP; they ~~can be~~ are for instance off-the-shelf computers or printers.

In some applications, only availability-critical devices need a double attachment, for instance the operator workplaces, while the majority of the devices are SANs. Taking advantage of the basic infrastructure of PRP, a DANP can be attached to two different bridges of the same LAN (e.g. a ring) and use protocols different from PRP to reconfigure the network in case of failure. The DANP then behaves as a bridge according to IEEE 802.1D. For instance, the bridge element may implement the RSTP protocol or a subset of RSTP where there is no forwarding of traffic between the ports. These abilities are optional and not detailed in this document. The supported mode is specified in ~~the PICS (see 5.9)~~ Clause 6.

4.1.6 Compatibility between singly and doubly attached nodes

Singly attached nodes (SAN), for instance maintenance laptops or printers that belong to one LAN, can be connected to any LAN. Bridges are always SANs. A SAN connected to one LAN cannot communicate directly to a SAN connected to the other LAN. These SANs are not aware of PRP redundancy, but DANPs generate a traffic that these SANs understand.

The condition is however that the SANs ignore the RCT in the frames, which ~~should be~~ is the case since a SAN cannot distinguish the RCT from the Ethernet padding.

Conversely, DANPs understand the traffic generated by SANs, since these do not append an RCT. They only forward one frame to their upper layers since the SAN traffic uses one LAN only.

4.1.7 Network management

A node has the same MAC address on both ports and has only one set of IP addresses assigned to that address. This makes redundancy transparent to the upper layers and especially allows ARP to work as with a SAN. Bridges in a LAN are not doubly attached devices, and therefore all managed bridges have different IP (and MAC) addresses. A

network management tool is preferably a DANP that can access nodes and bridges in both LANs.

4.1.8 Implication on ~~configuration~~ application

~~Since~~ When the same frame comes from the two ports with significant time difference, a frame could, in exceptional situations, be wrongly rejected as duplicate. ~~Therefore~~, there exists a mechanism to age out frames. This ageing mechanism causes in seldom cases that both duplicates are accepted ~~— a situation applications tolerate. Possibly the heaviest traffic comes from sampled data, whose period should be chosen considering the worst-case difference in propagation delays.~~ The duplicate must then be filtered out by the transport layer (e.g. TCP) or tolerated by the application. This is a normal case since a LAN does not guarantee absence of duplicates.

The ageing mechanism considers both the maximum repetition rate of the frames and the worst-case difference in propagation delay, as described in 4.1.10.3.

4.1.9 Transition to non-redundant networks

The mechanism of duplicate rejection can be implemented by a RedBox that does the transition between a SAN and the doubled LANs, as Figure 5 shows. The RedBox mimics the SANs connected behind it (called VDAN or virtual DANs) and multicasts supervision frames on their behalf. The RedBox is itself a DANP and has its own IP address for management purposes, but it may also perform application functions.




Figure 5 – PRP RedBox, transition from single to double LAN

4.1.10 Duplicate handling

4.1.10.1 Duplicate Accept mode (testing only)

The Duplicate Accept mode is used for testing purposes, to verify that duplicates are indeed discarded by the link layer and not by higher layer protocols.

In the Duplicate Accept mode, a sender is configured to send a normal frame over both frames without RCT ports.

In the Duplicate Accept mode, a receiver is configured to accept both frames and forward both of them (if both arrive) to its upper layers.

A node indicates its mode (Duplicate Accept/Duplicate Discard) in its supervision frames.

NOTE Upper layers also discard duplicates, but less effectively than the link layer.

4.1.10.2 Duplicate Discard mode

4.1.10.2.1 Principle

To allow the receivers to detect duplicates, the sender LRE appends a six-octet field that contains a sequence number, the Redundancy Control Trailer (RCT), to both frames it sends, as Figure 6 shows.

The receiver LRE uses the sequence number of the RCT and the source MAC address to detect duplicates. It forwards only the first frame of a pair to its upper layers. To offload the application processor, the LRE can be implemented with an independent pre-processor, an intelligent Ethernet controller or in hardware.

Keeping track of duplicates allows at the same time to improve the redundancy supervision. As an option, the receiver may remove the RCT before forwarding the (first) frame to its upper layer.

The RCT consists of the following fields:

- 16-bit sequence number (SeqNr);
- 4-bit LAN identifier (LanId);
- 12 bit frame size (LSDUsize);
- 16-bit suffix (PRPsuffix).




Figure 6 – PRP frame extended by an RCT

The LRE sends both (nearly identical) frames over both LANs.

The receiving LRE can then detect duplicates based on the RCT.

NOTE 1 This method considers that SANs also exist on the network, and that frames sent by SANs could be wrongly rejected as duplicates because they happen to have a trailing field with the correct PRP suffix, same sequence number and the same size. However, SANs send on one LAN only, therefore their source MAC address will not be the same as that of a frame on the other LAN, so that a frame from an SAN will never be discarded.

NOTE 2 Reception of a well-formed RCT is no proof that the source node is a DANP. Even if the path is wrong, a DANP nevertheless accepts the frame.

4.1.10.2.2 Use of PRPsuffix

The PRPsuffix identifies PRP frames and distinguishes PRP frames from other protocols that also append a trailer to their useful data.

4.1.10.2.3 Use of LSDUsize

The 12-bit LSDUsize field contains the length of the link service data unit (LSDU) in octets. The LSDUsize field allows the receiver to distinguish PRP frames from random, non-redundant frames as an additional check.

In VLANs, frame VLAN tags can be added or removed during transit through a bridge. To make the length field independent of VLAN tagging, only the LSDU and the RCT are considered in the LSDUsize, as Figure 7 shows.




Figure 7 – PRP VLAN-tagged frame extended by an RCT

If the receiver detects that the frame is closed by the correct PRP suffix, that the 12 bits before the PRP suffix correspond to the LSDU size, and that the LAN identifier matches the identifier of the LAN of the receiving port, the frame is a candidate for discarding.

Since short frames need padding to achieve the minimum frame size of 64 octets, the sender already includes the padding to avoid scanning the frame backwards, as Figure 8 shows. The minimum frame size is extended to 70 octets since the LRE or a RedBox can remove the RCT.




Figure 8 – PRP padded frame closed by an RCT

The padding size is increased to 74 octets when a VLAN tag is used.

NOTE A VLAN-tagged frame can pass several bridges, which removes or inserts VLAN tags. If the sender observes the ~~ISO/IEC 8802-3~~ Ethernet rule to send a minimum frame size of 68 octets for a VLAN-tagged frame and of 64 octets for a VLAN-untagged frame, there will never be a situation in which there is padding before and after the RCT.

4.1.10.2.4 Use of LanId

The 4 bit LanId field of the RCT carries a distinct identifier for LAN_A or LAN_B, specifically the codes 1010 ("A") and 1011 ("B"). Therefore, the A-frame and the B-frame differ in one bit (and in the FCS). The receiver checks that the frame comes from the correct LAN. It does not reject a frame that comes from the wrong LAN, since this could be a legitimate frame which

happens to have the PRP suffix and the length information in its last 28 bits, but it increments the error counters CntErrWrongLanA or CntErrWrongLanB since this hints at a configuration error. Since this kind of error is permanent, it is detected rapidly.

4.1.10.2.5 Use of SeqNr

The 16-bit field SeqNr is incremented by one for each frame a DANP sends. The doublet {source MAC address, sequence number} uniquely identifies copies of the same frame.

4.1.10.3 Duplicate discard algorithm and configuration

The algorithm for discarding duplicates is not specified. For instance, node tables, hash tables, fifos or sequence number tracking can be used. Whichever the algorithm, it is designed such that it never rejects a legitimate frame, while occasional acceptance of a duplicate can be tolerated. Since the 16-bit SeqNr wraps around after 65 536 frames and the sequence numbers are not necessarily contiguous, entries older than a specified time EntryForgetTime are purged. ~~This also allows node reboot at any sequence number position.~~

Figure 9 shows the conditions for discarding or accepting frames.

~~The time t_{wrapMin} is the minimum possible time between two repetitions of the same sequence number by legitimate frames after 65 536 increments (as dictated by the application and network constraints).~~

~~In a 100 Mbit/s network, t_{wrapMin} is about 400 ms at the theoretical maximum frame repetition rate from the same source.~~

~~The time $t_{\text{resideMax}}$ is the maximum time an entry remains in the duplicate table, it is smaller than EntryForgetTime.~~

~~Rules for reliable discard: $t_{\text{skew}} < t_{\text{resideMin}}$~~

~~Rule for safe accept: $t_{\text{residenceMax}} < t_{\text{aliasRepMin}}$~~

~~A node that reboots avoids generating aliases by not sending frames during a time longer than EntryForgetTime (see Table 4) on start-up. During this time, the receivers will age out the entry from the duplicate tables.~~

The time t_{skew} is the time difference between the arrivals of two copies of the same frame (from LAN A to LAN B and from LAN B to LAN A) in a particular case.

The maximum time t_{skewMax} between two copies is a network property, estimated by the network designer based on the number of bridges and the traffic for a particular application, e.g. 12 ms.

The time t_{aliasMin} is the minimum possible interval between two different, legitimate frames with the same sequence number (after 65 536 increments) in a particular application.

NOTE However, a node that reboots with an arbitrary sequence number could cause a shortening of time t_{aliasMin} , as will be detailed below.

In a 100 Mbit/s network, t_{aliasMin} is the theoretical minimum frame repetition interval of the smallest frame (64 octets) from the same source ($65\,536 \times 6,72\,\mu\text{s}$) = 440 ms.

In a 1 Gbit/s, t_{aliasMin} is 44 ms.

This is a worst case since applications rarely send minimum size frames contiguously.

EXAMPLE: An IEC 61850-9-2 device sending sampled values frames at 4,8 kHz will exhibit a $t_{aliasMin} = 13$ s.




Figure 9 – Duplicate Discard algorithm boundaries

Considering a device, entries in the duplicates list remain recorded during a time that varies between $t_{resideMin}$ and $t_{resideMax}$. These values depend on the implementation and are adequate for a number of applications, they can depend on the traffic if the list size is bound. It is recommended to make $t_{resideMin}$ and $t_{resideMax}$ adjustable depending on the data rate.

The following cases of Figure 9 are considered in Table 1:

Table 1 – Duplicate discard cases

Case	Condition	Description
Case 1	$t_{\text{skew}} > t_{\text{resideMin}}$	The duplicate is always discarded. This situation is desirable.
Case 2	$t_{\text{resideMin}} < t_{\text{skew1}} < t_{\text{resideMax}}$	The duplicate is sometimes not discarded if the entry was purged. This situation is tolerable if infrequent.
Case 3	$t_{\text{skew1}} > t_{\text{resideMax}}$	The duplicate is never discarded, the duplicate is treated as a new frame and creates a new entry in the duplicates list. This situation never occurs if the network was properly dimensioned with $(t_{\text{skewMax}} < t_{\text{resideMax}})$.
Case 4	$t_{\text{skew2}} > t_{\text{resideMax}}$	The duplicate is never discarded, but a legitimate frame is discarded since the entry created by LAN B conflicts with the alias on LAN A. This situation never occurs if the nodes are properly designed.
Case 5	$t_{\text{skew2}} > t_{\text{aliasMin}}$	The duplicate is discarded for the wrong reason, but this does not impair operation. This situation never occurs if the nodes are properly designed.

The conditions for correct operation are as follows:

- For the network engineer:

$$t_{\text{skewMax}} < (t_{\text{aliasMin}})/2 \text{ (no aliases)}$$

- For the device engineer:

$$t_{\text{resideMin}} > t_{\text{skewMax}} \text{ (recognizes all duplicates), and}$$

$$t_{\text{resideMax}} < (t_{\text{aliasMin}})/2 \text{ (prevents unwanted discards)}$$

In practice, $t_{\text{resideMin}}$ is often given by the number of frames kept in the buffer and therefore depends on the traffic, while $t_{\text{resideMax}}$ is enforced by a timer. This value is adjustable through the management setting $\text{IreDupListResideMaxTime}$ (see Clause 7)

A node that reboots could generate aliases at a lower interval than t_{aliasMin} since the sequence number at restart is arbitrary. To allow all possible destination nodes time to flush all entries of this node from the duplicates list, a node may not reboot faster than $\text{NodeRebootInterval}$ (see Table 8). $\text{NodeRebootInterval}$ holds for all nodes in the network and therefore it is independent from the link speed.

Therefore, all nodes flush their entries before $\text{NodeRebootInterval}$. The maximum residence time of an entry regardless of the speed (and of the value of $t_{\text{resideMax}}$) is EntryForgetTime .

The default value of EntryForgetTime in Table 8 (400 ms) allows in principle a t_{skewMax} of 800 ms, assuming a t_{aliasMin} of 1 600 ms. This fits all envisioned applications.

To comply with the above equations, $t_{\text{resideMax}}$ is usually much smaller than EntryForgetTime , especially in 1Gbit/s networks.

4.1.10.4 Frame size restriction

Appending the RCT could generate oversize frames that exceed the maxValidSize foreseen by ISO/IEC/IEEE 8802-3:2014 (1 522 octets).

Hence, it is expected that every network component in the LANs supports oversize frames up to maxValidSize foreseen by ISO/IEC/IEEE 8802-3:2014 plus 6 bytes of RCT (1 528 octets).

NOTE Many bridges are dimensioned for double-VLAN-tagged (non-ISO/IEC 8802-3 compliant Q-in-Q) frames that have a maximum size of 1 526 octets. Most Ethernet controllers are certified up to 1 528 octets. Most bridges would forward correctly frames of up to 1 536 octets.

4.1.11 Network supervision

The health status of each LAN and its attached devices (nodes and bridges) is monitored, otherwise redundancy helps little.

The receiver checks that frames are correctly received over both channels. It maintains error counters that network management can read.

To this effect, senders and receivers can maintain tables that record the last time a frame was received from another node, the time a multicast or broadcast frame was sent and other protocol information.

Supervision relies on each DANP sending periodically a PRP_Supervision frame that allows checking the integrity of the network and the presence of the nodes. At the same time, these frames allow checking which devices are DANP, the MAC addresses they use and which operating mode they support: Duplicate Accept or Duplicate Discard.

4.1.12 Redundancy management interface

Redundant devices and links are useless without network management supervising redundancy and calling for maintenance actions.

The LRE presents a network management interface that allows to track the health of each LAN, and especially to detect failures early when the error rate increases. To this effect, the LRE keeps for each port a counter of received messages and of messages received with an error.

The LAN statuses appear as SNMPv1 or SNMPv2/v3 objects, allowing the same tools to manage both nodes and bridges.

~~NOTE SNMP is part of the IP protocol suite.~~

4.2 PRP protocol specifications

4.2.1 Installation, configuration and repair guidelines

4.2.1.1 General

The following guidelines apply at installation time. They do not apply to conformance testing of the devices.

4.2.1.2 LANs layout

The network shall consist of two LANs that have similar properties, i.e. each one is able to carry the traffic that would exist in the absence of redundancy.

DANP shall be attached to both LANs.

SANs that need to communicate with one another shall be attached to the same LAN or to both LANs through a RedBox.

4.2.1.3 Labelling cables

The two LANs shall use cables distinctly identified as A and B.

If colours are used to identify the cables, A should be red and B should be blue.

NOTE The intention is that persons with colour disability distinguish them.

4.2.1.4 Labelling bridges

Bridges in the two LANs should have a distinct label or colour for each A or B.

4.2.1.5 Independent operation

The layout of both LANs shall fulfil the assumption of fail-independence.

4.2.1.6 Configuration

All DANPs shall be configured with the same LifeCheckInterval (see 4.5).

4.2.2 Unicast MAC addresses

Both ports A and B of a DANP shall be configured with the same MAC address.

This MAC address shall be unique in the network.

SANs connected to one LAN only shall not have the same MAC address as another node within the whole network (LAN_A plus LAN_B).

4.2.3 Multicast MAC addresses

All nodes in the network shall be configured to operate with the same multicast address for the purpose of network supervision (see 4.3).

4.2.4 IP addresses

The IP address(es) of any node or bridge shall be unique within the whole network (LAN_A plus LAN_B).

NOTE 1 ~~A device may have several~~ Uniqueness applies to all IP addresses of a device.

A DANP shall have the same IP address(es) when seen from either LAN_A or LAN_B.

Bridges on LAN_A and LAN_B are considered as SANs and shall have different IP addresses for the purpose of network management, even if they occupy the same location.

NOTE 2 A possible scheme is to distinguish IP addresses by one bit, e.g. allocate even IP addresses to LAN_A devices and odd IP addresses to LAN_B devices.

4.2.5 Nodes

4.2.5.1 Node types

Doubly attached nodes according to the parallel redundancy protocol (DANP) shall have two ports (portA and portB) that have the same abilities, and in particular could be used alternatively if only one LAN would be connected, portA being connected to LAN_A and portB to LAN_B.

Singly Attached Nodes (SAN) have only one port for the purpose of this protocol, no special requirements apply to them.

4.2.5.2 Labelling connectors

Subclause 4.2.5.2 applies to a DANP using two LANs of similar nature.

The connectors for each LAN shall be labelled distinctly as A and B.

When connectors are ordered vertically, LAN_A should be the upper connector and LAN_B the lower connector in the normal position of the device.

When connectors are ordered horizontally, the left connector should be the LAN_A and the right connector the LAN_B, as seen from the side where the cables or fibres are plugged.

The redundant connectors shall be independently removable and insertable.

4.2.6 Duplicate Accept mode (testing only)

4.2.6.1 Sending

A sender in the Duplicate Accept mode shall send the frame it receives from its upper layers unchanged over both its ports so that the two frames appear on the respective LANs.

4.2.6.2 Receiving

A receiver in the Duplicate Accept mode shall forward frames received from either port to its upper layers.

NOTE Duplicate Accept is a test setting not used for operation.

4.2.7 Duplicate Discard mode

4.2.7.1 Monitoring data set

A node shall maintain a monitoring data set consisting of the elements of Table 2.

Table 2 – Monitoring data set

Object name	Character	Description
SendSeq	mandatory	A 16-bit sequence counter used by this node for sending
CntErrorA and CntErrorB	optional	For each port A and B, a 32-bit counter indicating the number of erroneous frames received
CntReceivedA and CntReceivedB	optional	For each port A and B, a 32-bit counter indicating the number of frames received
CntErrWrongLanA and CntErrWrongLanB	optional	For each port A and B, a 32-bit counter indicating the number of mismatches
ActiveA and ActiveB	optional	For each port A and B, the status of the communication link
NodesTable	optional	A table of all other nodes that this node detects

4.2.7.2 NodesTable

A node may maintain a NodesTable for the purpose of network supervision. If a NodesTable is used, it shall consist of the elements given in Table 3 for each entry.

Table 3 – NodesTable attributes

Argument	Definition	Data Type
MacAddress	MAC address of the remote source node (6 octets)	OctetString6
CntReceivedA	Number of frames received from that node over LAN_A	Unsigned32
CntReceivedB	Number of frames received from that node over LAN_B	Unsigned32
CntErrWrongLanA	Number of frames that were received from that node with the wrong LAN identifier on LAN_A	Unsigned32
CntErrWrongLanB	Number of frames that were received from that node with the wrong LAN identifier on LAN_B	Unsigned32
TimeLastSeenA	Time at which the latest frame was received from that node over LAN_A	TimeTicks
TimeLastSeenB	Time at which the latest frame was received from that node over LAN_B	TimeTicks
SanA	True if the remote node is most probably a SAN accessible over portA	Boolean1
SanB	True if the remote node is most probably a SAN accessible over portB	Boolean1

A node shall clear a NodesTable entry of a given other node when the time elapsed since reception of a frame from that ~~a given~~ node over both TimeLastSeenA and TimeLastSeenB exceeds NodeForgetTime (see 4.5).

NOTE 1 The NodesTable is populated by the received frames and distinguishes SANs from DANPs. ~~A node could send frames without an RCT and over one port only if it can identify the destination as a SAN. The NodesTable allows a node to identify another node as SAN on one LAN only and in this case the node could send frames without an RCT over the port of that LAN only.~~ A node without a NodesTable always sends all frames with an RCT over both ports, also when the destination is a SAN. In applications where broadcast traffic is the bulk, the impact on network load is small.

NOTE 2 The key attribute of the NodesTable is MacAddress as received in the PRP_Supervision frame sent by a DANP or in any frame received from a SAN.

NOTE 3 Most of these attributes exist not only in one instance per physical remote node, but also as separate instances for each multi/broadcast address used by that node, and some also for each multi/broadcast address used by this (local) node.

NOTE 4 TimeTicks is a local time unit, not a global time.

4.2.7.3 Redundancy Control Trailer (RCT)

The Redundancy Control Trailer (RCT) appended to DANP frames shall consist of six octets, structured in the following way (in the order of transmission):

- a 16-bit sequence number (SeqNr) transmitted with the most significant 8 bits in the first octet, which reflects the counter SendSeq (see 4.2.7.1);
- a 4-bit LAN identifier (LanId) transmitted as the most significant 4 bits of the third octet, which carries the sequence “1010” for LAN_A and the sequence “1011” for LAN_B, respectively;
- a 12 bit LSDU size (LSDUsize) whose most significant 4 bits are transmitted in the least significant 4 bits of the third octet and the least significant 8 bits in the fourth octet, that indicates the size in octets of the LSDU starting from the end of the Protocol Type (PT) field as defined in ISO/IEC/IEEE 8802-3:2014 and IEEE 802.1Q (octet offset 12-13 without LAN header or 16-17 with VLAN header) to the RCT, excluding the PT and including possible padding and the RCT itself;
- a 16-bit PRPsuffix, set to 0x88FB.

NOTE 1 Padding inserted before the RCT is included in the LSDU size.

NOTE 2 The intention is that the LSDUsize is the same in PRP and HSR.

4.2.7.4 Sending (Duplicate Discard mode)

4.2.7.4.1 Sending

When sending a frame coming from its upper layers, a node shall:

If the destination node is not registered as a SAN in the NodesTable or if the node has no NodesTable	
	If the frame carries a IEEE 802.1Q tag and has a size smaller than 64 octets (without FCS)
	pad the frame to 64 octets as specified in ISO/IEC/IEEE 8802-3:2014
	If the frame does not carry a IEEE 802.1Q tag and has a size smaller than 60 octets (without FCS)
	pad the frame to 60 octets as specified in ISO/IEC/IEEE 8802-3:2014
	increment SendSeq (see 4.2.7.1), wrapping over through 0; append the 6-octet RCT with a LanId "A" for portA and LanId "B" for portB; append the 4-octet FCS; send the frame with LanId "A" through portA and the frame with LanId "B" through portB;
Else	
	send the frame to the port(s) where the SAN is registered;

NOTE 1 Simple nodes have no NodesTable and do not register remote nodes as DANP or SAN, they always append the RCT and send over both ports.

NOTE 2 The total padding to 70 octets (no VLAN tag) considers that the LRE or a RedBox could remove the 6-octet RCT. Padding to 74 octets (with VLAN tag) considers that in addition, an edge port in a bridge could remove the 4-octet priority tag. This way, the frame given to the destination respects the minimum frame size of 64 octets and the RCT is always just before the FCS.

4.2.7.4.2 Timing

The time skew between the sending of both frames of a pair should be kept at a minimum to reduce asymmetry of the transmission delays.

In particular, frames waiting for sending should be dropped from the sending queue if the corresponding port is not operational.

4.2.7.5 Receiver (Duplicate Discard mode)

4.2.7.5.1 Receiving

On reception of a frame over either port, the receiver shall:

If the frame is erroneous,	
	increment the error counter of the respective port CntErrorsA or CntErrorsB (see 4.2.7.1)
	ignore the frame;
Else if the destination MAC address belongs to the reserved addresses in IEEE 802.1Q-2004 Table 7-10,	
	treat the frame at the link layer;
Else if the frame is a duplicate (see 4.2.7.5.2 and 4.2.7.5.3),	
	discard the frame;
Else	
	forward the frame to the upper layers.

NOTE Frames with a destination MAC address according to IEEE 802.1Q:2004 2014, Table 8-1, Table 8-2 and Table 8-3 are treated at the link layer and are not expected to carry an RCT, even if sent by a DANP.

4.2.7.5.2 Identification of frames associated with the Duplicate Discard mode

A receiver shall identify a duplicate candidate as a valid frame:

- whose last 16 bits before the checksum match the 16-bit PRPsuffix,
- whose 12 bits before the PRPsuffix match the physical size of the LSDU (excluding the ProtocolType PT, but including the RCT) and
- whose LanId matches the LAN identity of the receiving port, 1010 (A) or 1011 (B).

4.2.7.5.3 Duplicate Discard method

The Duplicate Discard method shall not discard a frame sent as single or both frames of a pair, ~~while~~; exceptionally, both frames of a pair may be passed to the higher protocol layers.

The Duplicate Discard method shall forget an entry after EntryForgetTime (see Table 8).

4.2.7.5.4 Transparent reception

A receiver configured for TransparentReception (default) shall not remove the RCT before transferring the frame to its upper layers.

A receiver not configured for TransparentReception shall remove the RCT on frames where it has identified the presence of the RCT.

4.2.7.5.5 Receiving and NodesTable

If the node supports a NodesTable, it shall, at reception of a frame:

If this frame is not a PRP_Supervision frame or a HSR Supervision frame and its source MAC address is not yet in the NodesTable,	
	create an entry in the NodesTable for that source MAC address assuming it is a SanA or a SanB, depending which LAN the frame arrives on;
Else if this frame is a PRP_Supervision frame	
	If its source is not yet in the NodesTable
	create an entry in the NodesTable for that source assuming DANP Duplicate Accept or Duplicate Discard according to the PRP_Supervision frame contents;
	Else
	update the source status to DANP Duplicate Accept or Duplicate Discard;
record the local time at which the frame was received in the TimeLastSeenA, respectively TimeLastSeenB fields of the NodesTable for that source;	
increment by one (wrapping through 0) the counters CntReceivedA, respectively CntReceivedB of the NodesTable for that source.	
If it receives the frame from LAN_B from a node registered as SanA, or over LAN_A from a node registered as SanB,	
	set SanA = SanB = 1 for that source;
If the LanId does not match the identity of the port,	
	increment the CntErrWrongLanA, respective CntErrWrongLanB counter of the source device.

NOTE 1 Updating SanA and SanB allows moving an SAN from LAN_A to LAN_B and vice-versa. If this happens, the DANP will send on both LANs and after NodeForgetTime it will send only on the correct LAN.

4.3 PRP_Supervision frame

Each DANP shall multicast a PRP_Supervision frame over both its ports with the format specified in Table 4 (no VLAN) and Table 5 (with VLAN) respectively, every LifeCheckInterval (see 4.5).

Table 4 – PRP_Supervision frame with no VLAN tag

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0					msb	U/L	I/G								
2	PRP_DestinationAddress = multicast (01-15-4E-00-01-XX)														
4								lsb							
6					msb	U/L	0								
8	SourceAddress (MAC address of the node)														
10								lsb							
12	SupEtherType = 0x88FB														
14	SupPath				SupVersion < 64										
16	SupSequenceNumber														
18	TLV1.Type = 20 or 21							TLV1.Length = 6							
20					msb	U/L	0								
22	MacAddress (MAC address of the DANP)														
24								lsb							
26	TLV2.Type = 30							TLV2.Length = 6							
28					msb	U/L	0								
30	RedBoxMacAddress														
32								lsb							
34	TLV0.Type = 0							TLV0.Length = 0							
36															
Padding to 70 octets (no VLAN)															
60	SeqNr														
62	LanId (0x1010 or 0x1011)				LSDUsize = 52										
64	PRPsuffix = 0x88FB														
66	FCS														
68															
70															

Table 5 – PRP_Supervision frame with (optional) VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14																
16																
18																
20																
22																
24						msb	U/L	0								
26																
28									lsb							
30																
32						msb	U/L	0								
34																
36									lsb							
38																
40																
42																
44																
46																
48																
50																
52																
54																
56																
58																
60																
62																
64																
66																
68																
70																
72																
74																

NOTE 1 Octets with offset 14 to 17 are inserted only if a VLAN tag according to IEEE 802.1Q is used.

NOTE 2 The frame has a size of 70 octets (no VLAN tag) or 74 octets (with VLAN tag) to avoid padding if an entity removes the VLAN tag or the RCT.

NOTE 3 The SeqNr, LanId and LSDUsize are reused when the frame is passed to an HSR network. The SupEtherType is the same as the PRP suffix. It is also used for HSR.

NOTE 4 MAC addresses are assumed to have a length of 6 octets ~~in this example~~.

4.3.2 PRP_Supervision frame contents

The contents of a PRP_Supervision frame are defined in Table 6.

Table 6 – PRP_Supervision frame contents

Field name	Description
PRP_DestinationAddress	Multicast address 01-15-4E-00-01-XX reserved for the PRP protocol. By default XX is "00", but if conflicts arise, XX can be configured to take any value between 0x00 and 0xFF.
SourceAddress	MAC address of the node
SupEtherType	EtherType 0x88FB reserved for the PRP protocol.
SupPath	Reserved, set to 0
SupVersion	Protocol version, set to "1" for PRP. Implementation of version X of the protocol shall interpret version >X as if they were version X, ignoring any parameters and/or flags added by the more recent version, and interpret version <=X PRP_Supervision frames exactly as specified for the version concerned.
SupSequenceNumber	Sequence number incremented by 1 for each supervision frame sent.
TLV1.Type	Operation mode, a value of 20 indicates that the node supports the Duplicate Discard or a value of 21 indicates that it implements Duplicate Accept. Other values are reserved.
TLV1.Length	Length of the following MAC address in octets.
MacAddress	MAC addresses used by both port of the node.
TLV2.Type	Not present if the node is not a RedBox
TLV2.Length	Not present if the node is not a RedBox
RedBoxMacAddress	Not present if the node is not a RedBox
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

4.3.3 PRP_Supervision frame for RedBox

A RedBox, i.e. a node acting as a proxy for one or several SANs (called VDAN or virtual DAN) shall append to the TLV1 field a further TLV2 field with the content given in Table 7:

Table 7 – PRP_Supervision TLV for Redbox

Field name	Description
TLV1.Type	Operation mode with a value of 20 to indicate that the node supports the Duplicate Discard mode, since this is the only mode a RedBox supports. Other values are reserved.
TLV1.Length	Length of the following MAC address in octets.
MacAddress	MAC address of the VDAN
TLV2.Type	RedBox identifier with a value of 30 to indicate that the node is a RedBox
TLV2.Length	Length of the following MAC address in octets.
RedBoxMacAddress	MAC address of the RedBox that acts as proxy for the other device (VDAN)
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

4.3.4 Reception of a PRP_Supervision frame and NodesTable

When receiving a PRP_Supervision frame over any LAN, a node shall create an entry in the NodesTable corresponding to the MacAddress of that source as indicated in the message body, not in the source address, with the Duplicate Accept or Duplicate Discard mode as indicated in the frame.

If a node ceases to receive PRP_Supervision frames from a source for a time longer than NodeForgetTime, but receives frames from that source over one LAN only, it shall change the status of this node to SanA, respective SanB, depending on the LAN from which frames are received.

NOTE 1 This rule allows moving a SAN between LAN_A and LAN_B, and also to obtain the operating mode for a SAN if it was first registered at sending and not at receiving, since a DANP starts by sending on both LANs.

NOTE 2 This rule allows distinguishing an SAN from a DANP in Duplicate Accept mode with one line disconnected.

4.4 Bridging node

If the “bridging node” setting is enabled, the node shall act as a bridging node for its two ports, according to an unspecified protocol, i.e. HSR, RSTP or MRP, but shall not append an RCT to the frames it sends.

NOTE 1 The bridging node setting supports attachment of a DANP to two bridges of the same LAN to implement a partial redundancy topology. There is no requirement that normal frames ~~should~~ be bridged in case of a double failure, but implementers are free to include this feature.

~~NOTE 2 No RCT is appended when the bridging node setting is enabled.~~

NOTE 2 The presence of nodes in bridging mode is potentially dangerous for a PRP network since it breaches the assumption of fail-independence of the two LANs.

4.5 Constants

The constant parameters are shown in Table 8.

Other values may be defined at the user's responsibility.

Table 8 – PRP constants

Constant	Description	Default value
LifeCheckInterval	~~How often a node sends a PRP_Supervision frame~~ Typical interval between two successive PRP_Supervision frames sent by the same node.	2 000 ms
NodeForgetTime	Time after which a node entry is cleared from the NodesTable after frames from this node cease to be received.	60 000 ms
EntryForgetTime	~~Time after which an entry is removed from the duplicate table~~ Maximum time an entry may reside in the duplicate table, independently from speed.	400 ms
NodeRebootInterval	Minimum time during which a node that reboots remains silent, independently from speed.	500 ms

4.6 PRP service specification

Objects in the LRE, e.g. error counters, can be accessed from an application running on the same node using direct calls or from an application on another node using a management protocol such as SNMP or a specific protocol such as IEC 61850.

All LRE objects are described in Clause 7 as a Management Information Base (MIB) conformant to SNMP (RFC 2578, RFC 3418), which is common for PRP and HSR.

5 High-availability Seamless Redundancy (HSR)

5.1 HSR objectives

Clause 5 describes the application of the PRP principles of Clause 4 to implement a High-availability Seamless Redundancy (HSR), retaining the PRP property of zero recovery time, applicable to any topology, in particular rings and rings of rings.

NOTE With respect to PRP, HSR allows to roughly halve the network infrastructure. With respect to rings based on IEEE 802.1D (RSTP), IEC 62439-2 (MRP), IEC 62439-6 (DRP) or IEC 62439-7 (RRP), the ~~available network bandwidth for network traffic is somewhat reduced depending on the type of traffic~~ network traffic is about the same.


Nodes within the ring are restricted to be HSR-capable bridging nodes, thus avoiding the use of dedicated bridges. Singly Attached Nodes (SANs) such as laptops or printers cannot be attached directly to the ring, but need attachment through a Redundancy Box (RedBox).

5.2 HSR principle of operation

5.2.1 Basic operation with a ring topology

As in PRP, a node has two ports operated in parallel; it is a DANH (Doubly Attached Node with HSR protocol).

A simple HSR network consists of doubly attached bridging nodes, each having two ring ports, interconnected by full-duplex links, as shown in the example of Figure 10 (multicast) and Figure 11 (unicast) for a ring topology.



IEC


Figure 10 – HSR example of ring configuration for multicast traffic

A source DANH sends a frame passed from its upper layers ("C" frame), prefixes it by an HSR tag to identify frame duplicates and sends the frame over each port ("A"-frame and "B"-frame).

A destination DANH receives, in the fault-free state, two identical frames from each port within a certain interval, removes the HSR tag of the first frame before passing it to its upper layers ("D"-frame) and discards any duplicate.

The nodes support the IEEE 802.1D bridge functionality and forward frames from one port to the other, ~~except if they already sent the same frame in that same direction~~ according to four rules:

- 1) a node will not forward a frame that it injected into the ring or mesh;
- 2) a node will not forward a frame for which it is the unique destination (except for special applications such as redundancy supervision);
- 3) a port will not send a frame that is a duplicate of a frame that it already sent into that same direction;
- 4) a port will (optionally) refrain from sending a frame that is a duplicate of a frame that it already received from the opposite direction (except for supervision and timing frames).



IEC

Figure 11 – HSR example of ring configuration for unicast traffic

~~A destination node of a unicast frame does not forward a frame for which it is the only destination, except for testing.~~

Frames circulating in the ring carry the HSR tag inserted by the source, which contains a sequence number. The doublet {source MAC address, sequence number} uniquely identifies copies of the same frame.

NOTE 1 The maximum time skew between two frames of a pair depends on the relative position of the receiving node and of the sending node. Assuming a worst case in which each node in the ring is transmitting at the same time its own frame with the largest size of 1 536 octets (maximum length supported by the EtherType defined in ISO/IEC/IEEE 802.2 8802-2:1998), each node could introduce 125 µs of delay at 100 Mbit/s. With 50 nodes, the time skew may can exceed 6 ms.

NOTE 2 The minimum delay before a frame returns to the node that injected it depends on the ring size. Assuming a worst case in which there is no other traffic and in which each node in the ring operates in cut-through with 4 μ s residence time, in a ring of 32 nodes, a frame with the maximum size of 1 536 octets (125 μ s at 100 Mbit/s) returns to the injecting node after 124 μ s, i.e. before the injecting node finished transmitting it.

5.2.2 DANH node structure

Figure 12 shows a conceptual view of the structure of a DANH implemented in hardware; practical implementations can be different. The two HSR ports A and B and the device port C are connected by the LRE, which includes a switching matrix allowing to forward frames from one port to the other. The switching matrix allows cut-through bridging. The LRE presents to the higher layers the same interface as a standard Ethernet transceiver would do.

The input circuit checks if this node is the destination of the frame and possibly does VLAN and multicast filtering to offload the processor. The Duplicate Discard is implemented in the output queues.




Figure 12 – HSR structure of a DANH

5.2.3 Topology


5.2.3.1 Attachment of singly attached nodes

Singly attached nodes (SAN), for instance maintenance laptops or printers cannot be inserted directly into the ring since they have only one port and cannot interpret the HSR tag in the frames. SANs communicate with ring devices through a RedBox that acts as a proxy for the SANs attached to it, as shown in Figure 10 and Figure 11. The RedBox is detailed in 5.2.3.8.

Connecting non-HSR nodes to ring ports, breaking the ring, is allowed to enable configuration. Non-HSR traffic within the closed ring is supported in an optional mode.

5.2.3.2 Use of HSR with separate LANs

HSR nodes can be connected in the same way as PRP nodes. To this effect, an HSR node may be set to “no forwarding” (mode N). Unlike PRP, SANs cannot be attached directly to such a duplicated network unless they are able to interpret the HSR tag. ~~The details of this operation mode are not covered in this version of the standard.~~



IEC

Figure 13 – HSR example of topology using two independent networks

5.2.3.3 Peer coupling of rings

Two HSR rings may be connected by quadruple port devices with forwarding capabilities, called QuadBoxes, as Figure 14 shows. This is advantageous when the traffic flow exceeds the capabilities of a single ring. However, transmission delays from end to end are not improved.

Although one QuadBox is sufficient to conduct the traffic in the fault-free state of the network, two QuadBoxes are used to prevent a single point of failure.

A Quadbox forwards frames over each ring as any HSR node, and passes the frames unchanged to the other ring, except if the frame can be identified as a frame not to be forwarded to the other ring. To this effect, a QuadBox is expected to filter traffic based for instance on multicast filtering or on VLAN filtering. There is no learning of MAC addresses in a QuadBox, ~~though~~, since the learning of MAC addresses on specific ports of a QuadBox device could lead to a short break in communication, if the QuadBox that has learned an address and is forwarding network traffic, fails.

With QuadBoxes realized as single physical entities, the two interconnected rings share the same redundancy domain concerning fault tolerance. If one QuadBox breaks down, both interconnected rings are in a degraded state and cannot tolerate a further fault.

Therefore, constructing QuadBoxes in the same way as a RedBox can help keep the redundancy independent. The QuadBox then consists of two devices connected by an interlink. For this reason, the RedBox specifications include the HSR connection.




Figure 14 – HSR example of peer coupling of two rings

The presence of two QuadBoxes on the same ring causes that two copies of the same frame are transferred from the first ring to the second, each generating another two copies.


This does not cause four frames to circulate on the second ring, since, when a copy from a first QuadBox reaches the second QuadBox on the same second ring, the second QuadBox will not forward it if it already sent a copy that came from its interlink. Conversely, if the second QuadBox did not yet receive a copy from its interlink, it will forward the frame, but not the copy that comes later from the interlink.

When a QuadBox receives a frame that it itself injected into the ring or a frame that the other QuadBox inserted into the ring, it forwards it to the interlink and to its other port if it did not already sent a copy. This duplicate will be discarded at the other end of the interlink. This scheme may cause some additional traffic on the interlink, but it allows to simplify the design of the logic.

NOTE The maximum time skew between two frames of a pair is about the same as if all nodes were on the same ring.

5.2.3.4 Hierarchical ring topology

An HSR network may consist of rings connected by QuadBoxes, as Figure 15 shows.



IEC

Figure 15 – HSR example of connected rings


Although a single QuadBox is sufficient to sustain the traffic, two independent QuadBoxes are needed to avoid a single point of failure.

Some SANs are connected directly to the DANH that performs the duty of a simplified RedBox.

5.2.3.5 Connection of an HSR ring to a PRP network

An HSR may be coupled to a PRP network through two RedBoxes, one for each LAN, as Figure 16 shows. In this case, the RedBoxes are configured to support PRP traffic on the interlink and HSR traffic on the ring ports.

The sequence number from the PRP RCT is reused for the HSR tag and vice versa, to allow frame identification from one network to the other and to identify pairs and duplicates on the HSR ring, introduced by a twofold injection into the ring through the two HSR RedBoxes.



IEC

Figure 16 – HSR example of coupling two redundant PRP LANs to a ring

The HSR RedBoxes for connecting the ring to a PRP network operate identically to the HSR RedBoxes used to attach SANs described in 5.2.3.1, except that they are configured as RedBox “A” or RedBox “B” to accept PRP frames on their interlink.


In Figure 16, RedBox A and RedBox B would send the same frame (A and AB and B and BA, respectively), but if a RedBox receives the frame before it could send it itself, it refrains from sending it.

In the example of Figure 16, RedBox A will not generate an “A” frame on behalf of LAN A if it previously received the same frame as “AB” from the ring, or conversely, RedBox “B” will generate an “AB” frame if it did not previously receive an “A” frame from the ring, which is the case whenever frame “A” is not a multicast frame.

Multicast frames or unicast frames without a receiver in the ring (void arrows in Figure 16) are removed by the RedBox that inserted them into the ring, if they originated from outside the

ring. To this purpose, the frames carry a LAN identifier that is at the same time the Redbox identifier.

Figure 17 shows the same coupling when the source is within the ring.



IEC


Figure 17 – HSR example of coupling from a ring node to redundant PRP LANs

To avoid reinjection of a frame into the PRP network through the other RedBox, each HSR frame carries the identifier of the PRP network from which the frame came originally, the 4-bit PathId. Therefore, RedBoxes are configured and identified by the NetId PathId of the PRP network LAN to which they are attached.

5.2.3.6 Connection of several PRP networks to an HSR ring

Up to 7 PRP networks can be connected to the same HSR ring, each being identified by a 3-bit NetId.

The two RedBoxes that connect a PRP network with an HSR ring are configured with the NetId (1 to 7) and the LanId (A=0/B=1), which form the 4-bit PathId (see Figure 18).



IEC

Figure 18 – HSR example of coupling from a ring to two PRP LANs


To prevent reinjection of frames coming from one PRP network into another PRP network or from the other LAN of the same PRP network, a RedBox only forwards from the HSR ring frames that do not carry its own NetId PathId.

When inserting into the ring a PRP frame from LAN A or from LAN B of a PRP network with a given NetId, a RedBox inserts into the PathId of the HSR tag its own NetId and the LanId, i.e. one of "2"/"3", "4"/"5", "6"/"7", "8"/"9", "A"/"B", "C"/"D" or "E"/"F", depending if it is RedBox A or RedBox B.

Conversely, when forwarding a frame from the ring to a PRP network, a RedBox inserts the LanId "A" or "B" into the RCT, depending if it is RedBox A or RedBox B.

5.2.3.7 Connection of one PRP networks to several HSR rings

A PRP network can be connected to any number of HSR rings, but these rings cannot be connected between themselves, neither by QuadBoxes nor by another PRP network since this would create loops (see Figure 19).




IEC

Figure 19 – HSR example of coupling three rings to one PRP LAN

5.2.3.8 Meshed topology

HSR allows any kind of meshing, and provides redundancy as long as the structure is free from single point of failure. For instance, Figure 20 shows a matrix arrangement of nodes. In this case, nodes have more than two ports operated in parallel that operate like the QuadBoxes. A frame received from one port is forwarded to all other ports except the one that received it, and each port forwards the frame unless it already sent a duplicate.




IEC

Figure 20 – HSR example of meshed topology

5.2.4 RedBox structure

Figure 21 shows the general structure of a RedBox:



IEC

Figure 21 – HSR structure of a RedBox

The RedBox has an LRE that performs the duties of the HSR protocol, in particular:

- forwards the frames received from one HSR port to the other HSR port, unless the frame was sent already;
- receives frames addressed to its own upper protocols;
- prefixes the frames sent by its own upper layers with the corresponding HSR tag before sending two copies over its HSR ports.

A RedBox can be operated in one of three modes: as a SAN connection, as a PRP connection or as an HSR connection (half of a QuadBox). Depending on the mode of operation, the frame handling at the interlink interface of the RedBox differs (see 5.4.2 and 5.4.4).

The switching logic in Figure 21 can be incorporated into the RedBox, so the interlink becomes an internal connection.

A simple RedBox is present in every node, since the LRE makes a transition to a single non-HSR host. Also, it is usual to have more than one host in a node, since a port for maintenance often exists.

5.3 HSR node specifications

5.3.1 HSR operation

5.3.1.1 Node properties

A DANH shall be operable in one of the following modes, which shall be changeable at runtime using management commands:

- Mode H (mandatory, default mode): HSR-tagged forwarding – in this mode, the DANH shall insert the HSR tag on behalf of its host and shall forward the ring traffic, except for frames sent by the node itself, duplicate frames and frames for which the node is the ~~unicast~~ **unique** destination ~~shall not be forwarded~~.
- Mode N (optional): No forwarding – in this mode, the node shall behave like mode H with the exception that it shall not forward ring traffic from port to port.
- Mode T (optional): Transparent forwarding – in this mode the DANH shall remove the HSR tag before forwarding the frame to the other port and **shall** send a frame from the host to both ports, untagged and without discarding duplicates.
- Mode M (optional): Mixed forwarding HSR-tagged and non HSR-tagged – in this mode, the DANH shall insert the HSR tag depending on local criteria when injecting frames into the ring. HSR tagged frames from ring ports are handled according to Mode H. Non-HSR tagged frames from ring ports are handled according to IEEE 802.1D forwarding rules.
- Mode U (optional): Unicast forwarding – in this mode, the node shall behave as in Mode H, except that it shall **also** forward ~~unicast~~ traffic for which it is the **unique** destination ~~as it does for multicast~~.
- Mode X (optional): No sending on counter-duplicate – in this mode, the node shall behave as in Mode H, except that a port shall not send a frame that is a duplicate of a frame that it received completely and correctly from the opposite direction.

5.3.1.2 Host sequence number

An HSR node shall maintain a sequence number for each MAC address the host uses. The sequence number is initialized with 0.

NOTE The LRE is expected to detect the MAC address of the host by listening to the frames it receives from it. The host MAC address can also be configured into the LRE. Without knowledge of the host MAC address, the LRE forwards all HSR traffic to the host, treating unicast frames as multicast.

5.3.1.3 Host multicast address

A node shall have the ability to filter out at least MulticastFilterSize (see Table 11) multicast addresses or multicast address groups. If the table is not configured, the LRE shall forward all multicast frames to the host and to the other port, duplicates excluded.

5.3.1.4 Host unicast address

A node shall have the ability to communicate its MAC address(es) to the LRE. If the MAC address is not configured, the LRE shall forward all unicast frames to the host and to the other port, duplicates excluded.

5.3.1.5 Host reboot

When a node reboots, it shall refrain from sending frames over its HSR ports during a time NodeRebootTime (see Table 11) and preferably start transmission by sending a supervision frame.

5.3.2 DANH receiving from its link layer interface

For each frame to send on behalf of its link layer interface, a source node shall:

If the node is in Mode T		
	do not modify the frame;	
	duplicate the frame, enqueue it for sending into both HSR ports	
Else if the node is in Mode M		
	If the local criteria is not met	
		insert the HSR tag and increment the sequence number depending on local criteria; duplicate the frame, enqueue it for sending into both HSR ports
	Else	
		do not insert the HSR tag; process the frame according to IEEE 802.1D rules
Else		
	insert the HSR tag with the sequence number of the host; increment the sequence number, wrapping through 0; duplicate the frame, enqueue it for sending into both HSR ports	

NOTE 1 Enqueuing means that the frame will be sent as soon as no former or higher-priority frames are in the queue and the medium is ready.

NOTE 2 Sending a non-HSR frame to both ports ~~should~~ does not ~~cause~~ produce circulating frames since such frames will not be forwarded by the adjacent HSR node. This mechanism is intended for off-ring configuration of an HSR node through a normal PC.

5.3.3 DANH receiving from an HSR port

A node receiving a frame from one of its HSR ports shall:

If this frame is not HSR-tagged:		
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)	
		consume the frame and do not forward it.
	Else if the node is in Mode M	
		forward according to 802.1D forwarding rules
	Else if the node is in Mode T	
		enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.
	Else	
		enqueue the unchanged frame for passing to its link layer interface. do not forward this frame
Else (HSR-tagged frame):		
	If the node is in Mode T	
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface
	Else (node is not in Mode T)	
	If this node is a destination:	
		If this is the first occurrence of the frame over the link layer interface:
		remove the HSR tag and pass the modified frame to its link layer interface.
		Else (this is not the first occurrence of the frame over the link layer interface):
		do not pass the frame to the link layer interface.
	Else (if this node is not a destination):	
		do not pass the frame to the link layer interface.
	If this node is not the only destination ~~(multicast or unicast) for another node~~ or the node is in Mode U:	
		If this is the first occurrence of the frame and the node is not in Mode N:
		enqueue the unmodified frame for sending over the ~~second~~ other port(s)
		Else (this is not the first occurrence of the frame):
		discard the frame.
	Else (If this node is not in Mode U and it is the ~~only (unicast)~~ unique destination):	
		discard the frame.

NOTE 1 It is possible that more than one duplicate arrives, especially when rings are coupled.

NOTE 2 A node accepts an HSR tagged frame also if the LanId does not correspond to the PortId and if the LSDUsize does not match the frame size.

5.3.4 DANH forwarding rules

A node shall not send over a port, a frame that is a duplicate of a frame previously sent over that port in that same direction. ~~This can also be seen in the behavioural description in 5.3.3.~~

A node that receives a frame shall not modify its source address or HSR tag when forwarding the frame even if it did modify the contents of the frame.

A node that detects on the base of the signal quality supervision or FCS that the frame is damaged or truncated, shall not forward it. However, if the node operating in cut-through already started forwarding and then detects that the frame is damaged or truncated, it shall append the error sequence foreseen in ISO/IEC/IEEE 8802-3:2000 2014, 27.3.1.2.2 and then stop transmission of that frame.

If a previously connected port is disconnected from the network, a node shall purge the port's buffer so that it cannot send an obsolete frame, and resume buffering when the port is reconnected.

For frames that originate from a DANH, LanId may be set to 0 or to 1 to indicate the sending port (A = 0, B = 1)

Cleaning the transmit queues from entries coming from this node is not necessary, since a transmit queue empties itself whether the attached link is operational or not.

If a port received a complete and correct frame, it could in principle refrain from sending a duplicate of that frame in the opposite direction, since it can be assumed that all nodes in that direction already received correctly the frame since they forwarded it. This method saves bandwidth, but there are several caveats:

- This method is not applicable to supervisory frames, to preserve integrity check of the ring in both directions.
- This method is not applicable to PTP frames (see Annex A) since the delay is different in both directions.
- This method is not applicable when Mode "U" is enabled, to allow all nodes to receive all traffic.
- This method is not effective for the traffic that is not forwarded by a unique destination.
- This method only reduces the broadcast and multicast traffic when the node is not operating in cut-through mode.
- This method introduces a variable response delay depending on the position of the node in the ring and on race conditions; it cannot be predicted which node, if any, removes the duplicate.
- This method does not guarantee that the gained bandwidth remains available for other traffic.
- This method does not improve the worst case network delay, when a frame or a node is lost. Nodes presenting this behaviour would interwork with nodes not implementing it.

The ability to refrain from sending on counter-duplicate is expressed as Mode_X in the PICS.

NOTE 1 These rules remove circulating HSR frames and open the ring, in the same way as an RSTP or similar protocol. It applies to frames originally sourced by the node and to frames circulating in case a device is removed after having sent a frame, and the ring is closed again, for instance by a mechanical bridging device or when a DANH is powered down. In a ring of 50 nodes, a frame comes back to its originator after some 6 ms.

NOTE 2 These conditions enable a node to operate either in store-and-forward or in cut-through mode. Delaying the forwarding of a frame does not affect the worst-case ring delay.

NOTE 3 The Duplicate Discard method of PRP is not a preferred method for discarding duplicates in HSR, since HSR aims at preventing duplicates from circulating.

NOTE 4 The fact that the sequence numbers of the frames sent by one source are not monotonically increasing is not a reason for discarding the frame. This observation can however be used for supervision of the network.

NOTE 5 In cut-through operation, the complete reception of the HSR tag takes approximately 2,4 µs at 100 Mbit/s, after which the node decides to forward the frame or not. In store-and-forward, the complete reception of a frame takes some 123 µs at 100 Mbit/s for the maximum size frame (1 522 octets).

~~NOTE 6 Cleaning the transmit queues from entries coming from this node should not be necessary, since a transmit queue should empty itself whether the attached link is operational or not.~~

5.3.5 CoS

For the operation of HSR, priorities and VLANs are not required.

An HSR node is expected, as expressed in its PICS:

- to support at least 2 levels of priority according to IEEE 802.1D (IEEE 802.1p);
- to filter VLAN traffic according to IEEE 802.1Q;
- to filter multicast traffic.

5.3.6 Clock synchronization

Handling of clocks as per IEC 61588 is described in Annex A.

5.3.7 Deterministic medium access

Deterministic medium access is ~~described in Annex B~~ not specified.

NOTE Deterministic medium access can take advantage of the precise clock synchronization, by letting all nodes start transmission of frames with time-critical data at the same time in a reserved time slot and ensuring that all other frames completed transmission before the start of that time slot. The size of a frame is announced in the HSR header.

5.4 HSR RedBox specifications

5.4.1 RedBox properties

A RedBox is a device with at least three ports, two of them being ring ports for the HSR protocol, the third port being connected to an interlink.

A RedBox shall behave as a DANH for all traffic for which it is the source or the destination.

In addition to the DANH behaviour, the behaviour of the RedBox Interlink is described in 5.4.2. Since the RedBox distinguishes, on reception from the HSR network, between its internal interfaces and its interlink, the behaviour of a RedBox upon reception from an HSR network is described in 5.4.4.

A RedBox shall have a configurable NetId that is composed of three bits as defined in 5.7.1. It is used in the HSR tag for frames that were received from the interlink port of the RedBox. An HSR RedBox connection to a PRP LAN will not forward HSR frames of the configured NetId into the PRP LAN.

The NetId of the RedBox shall not be used in the HSR tag of frames generated within the RedBox. Such frame shall use a NetId of zero in the HSR tag instead, or a ring NetId that can be configured separately, with a default value of zero.

A RedBox shall be configurable for one of three modes:

- 1) HSR-SAN: the traffic on the interlink is not HSR, not PRP;
- 2) HSR-PRP: the traffic on the interlink is PRP-tagged as “A” or “B”;
- 3) HSR-HSR the traffic on the interlink is HSR-tagged.

NOTE 1 A RedBox itself implements at least a management interface.

NOTE 2 A RedBox is expected to have its own IP address, especially for configuration messages. It can be accessed over the interlink or over the HSR ports.

NOTE 3 The interlink can be an internal connection if the RedBox serves as bridge at the same time.

5.4.2 RedBox receiving from interlink

When receiving a frame from its interlink port, a RedBox shall:

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

If the frame carries a HSR tag:			
	If the node is in Mode T		
		remove the HSR tag enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else if the RedBox operates in HSR-HSR mode		
		If the RedBox is a destination of the frame	
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else (If this is the first occurrence of the frame at the link layer interface)
			remove the HSR tag enqueue to the link layer interface of the RedBox
		If the frame is to be injected into the ring (RedBox is not only destination, Multicast/VLAN is ok)	
			If this is not the first occurrence of the frame
			discard the frame (already sent over that port)
			Else (If this is the first occurrence of the frame)
			enqueue the unmodified frame into each HSR port
	Else (If the RedBox does not operate in HSR-HSR mode)		
		discard the frame	
Else if the frame carries a PRP RCT and the RedBox is in HSR-PRP mode			
	If the node is in Mode T		
		remove the PRP RCT enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else		
		If the source MAC address is not already registered:	
			create an entry in the ProxyNodeTable;
		If the PRP tag does not correspond to the LanId of the RedBox "A" or "B"	
			increment the error counter lreCntWrongLanIdC;
		If the RedBox is a destination of the frame	
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else
			enqueue to the link layer interface of the RedBox (with the PRP RCT)
		If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok)	
			If this is not the first occurrence of the frame at each HSR port
			discard the frame (already sent over that port)
			Else (If this is the first occurrence of the frame at each HSR port)
			reuse the PRP sequence number and the RedBox's path identifier to build the HSR tag enqueue the frame into each HSR port
Else (if the frame carries neither a HSR tag nor a PRP RCT)			
	If the node is in Mode T		

		enqueue the unmodified frame to the link layer interface of the RedBox enqueue the unmodified frame into each HSR port
	Else if the node is in Mode M	
		forward according to 802.1D forwarding rules
	Else	
		If the source MAC address is not already registered in the ProxyNodeTable:
		create an entry in the ProxyNodeTable;
		Else (If the source is already registered)
		(void)
		If the RedBox is a destination of the frame:
		enqueue to the link layer interface of the RedBox
		If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok))
		append the HSR tag using the sequence number of that node; increment the sequence number of that source; enqueue the tagged frame into each HSR port

NOTE 1 Reception of an HSR frame over the interlink is considered as a configuration error in PRP connection mode. If the RedBox is used as one half of a QuadBox in HSR connection mode, then the interlink will only carry HSR traffic. In HSR-HSR mode, the interlink is handled as an additional HSR port intended for QuadBox coupling and thus does not use a ProxyNodeTable, unlike PRP-HSR and SAN-HSR RedBoxes. The RedBox implements a NodesTable to record the presence of HSR nodes.

NOTE 2 The RCT is only removed if the DANH identifies positively the sender as a DANP if the node has a NodeTable. Nodes without a NodeTable remove the RCT according to the mode setting.

5.4.3 RedBox forwarding on the ring

In addition to the forwarding rules of 5.3.4, a RedBox shall not forward in the ring a unicast frame that is intended for one of the nodes that are registered in the ProxyNodeTable. This condition is enabled by default and can be disabled by setting the RedBox to Mode U.

5.4.4 RedBox receiving from an HSR port

A RedBox that receives a valid frame over one HSR port shall:

If this frame is not HSR-tagged:			
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)		
		consume the frame and do not forward it.	
	Else if the node is in Mode M		
		enqueue the unchanged frame for passing to the interlink. forward according to 802.1D forwarding rules	
	Else if the node is in Mode T		
		enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.	
	Else		
		enqueue the unchanged frame for passing to its link layer interface. do not forward this frame	
Else (frame is HSR-tagged):			
	If the node is in Mode T		
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface	
	Else		
		If this node is a destination:	
			If this is the first occurrence of the frame over the link layer interface:
			remove the HSR tag and pass the modified frame to its link layer interface.
			Else (this is not the first occurrence of the frame over the link layer interface):
			do not pass the frame to the link layer interface.
	Else (if this node is not a destination):		
			do not pass the frame to the link layer interface.
		If this node is not the ~~only~~ unique destination ~~(multicast or unicast for another node)~~ or the node is in Mode U:	
			If this is the first occurrence of the frame and the node is not in Mode N:
			enqueue the unmodified frame for sending over the ~~second~~ other port(s).
			Else (this is not the first occurrence of the frame or the node is in Mode N):
			do not enqueue the unchanged frame to the second HSR port.
			If this is the first occurrence of the frame in direction of the interlink:
			If the RedBox is in SAN mode:
			remove the HSR tag;
			If Source in the ProxyNodeTable
			discard the frame
			Else
			enqueue frame for passing to the interlink.
			Else if the RedBox is in PRP mode:
			If the netId matches that of the RedBox
			discard the frame
			Else

					remove the HSR tag and append the PRP RCT with LanId ("A" or "B") of the RedBox and reusing the HSR sequence number.
					enqueue frame for passing to the interlink.
				Else (if the RedBox is in HSR mode)	
					do not modify the frame.
					enqueue frame for passing to the interlink.
				Else (If this is not the first occurrence of the frame in direction of the interlink):	
					discard the frame
				Else (If this node is the only (unicast) destination) and the node is not in Mode U:	
					discard the frame.

NOTE For HSR-HSR connection only, a RedBox does not check if the frame was sent by one of the nodes for which it is a proxy since it cannot distinguish if the frame could have been sent by a redundant RedBox.

5.4.5 RedBox receiving from its link layer interface

A RedBox that receives a valid frame from its link layer interface shall:

If the node is in Mode T	
	do not modify the frame;
	duplicate the frame, enqueue it for sending into both HSR ports
Else If the node is in Mode M	
	insert the HSR tag and increment the sequence number depending on local criteria forward depending on local criteria
Else	
	If the RedBox is in SAN mode:
	enqueue frame for passing to the interlink.
	Else If the RedBox is in PRP mode:
	insert the PRP RCT with the sequence number of the host; enqueue frame for passing to the interlink.
	Else (if the RedBox is in HSR mode)
	insert the HSR tag with the sequence number of the host; enqueue frame for passing to the interlink.
	duplicate the frame, enqueue it for sending into both HSR ports

5.4.6 Redbox ProxyNodeTable handling

A RedBox shall hold a ProxyNodeTable containing an entry for each represented association, and supporting at least ProxyNodeTableMaxEntries entries.

A RedBox shall purge the entry of a node in the ProxyNodeTable with a configurable time-out (default of ProxyNodeTableForgetTime) for non-receiving frames of this node.

NOTE The actual size of the ProxyNodeTable is indicated in the PICS.

5.4.7 RedBox CoS

Same as 5.3.5, except that the RedBox is expected to support more VLANs and more complex and comprehensive options for engineering traffic flow in the network on higher protocol layers, e.g. multicast filtering rules, than a simple node.

5.4.8 RedBox clock synchronization

Handling of clocks as per IEC 61588 is described in Annex A.

5.4.9 RedBox medium access

Same as 5.3.7, except that the RedBox shall be able to aggregate several nodes for which it acts as a proxy.

5.5 QuadBox specification

A Quadbox shall operate conceptually as two RedBoxes in HSR-HSR mode, back-to-back.

5.6 Duplicate Discard method

For the purpose of Duplicate Discard, a frame shall be identified by:

- its source MAC address;
- its sequence number.

NOTE In addition, it is possible to use other fields of the frame such as the checksum to aid in duplicate detection.

The Duplicate Discard method is not specified.

Any Duplicate Discard method shall be able to forget an entry identified by <Source MAC Address><Sequence number> after a time EntryForgetTime.


5.7 Frame format for HSR

5.7.1 Frame format for all frames

HSR frames are identified uniquely by their HSR tag.

The HSR tag consists of:

- 16-bit Ethertype (HSR_EtherType = 0x892F)
- 4-bit path identifier (PathId)
- 12 bit frame size (LSDUsize)
- 16-bit sequence number (SeqNr)



IEC

Figure 22 – HSR frame without a VLAN tag

The 4-bit PathId field prevents reinjection of frames coming from one PRP network to another PRP network.

The 3 most significant bits of the 4-bit path field are the NetId with values 0 for regular HSR frames, 1 to 6 for frames originating from a PRP network (up to 6 networks supported, might also be used for other HSR networks in the future), 7 is reserved.

The least significant bit of the 4-bit path field is the ~~LaneId~~ **LanId**. This bit indicates for a frame coming from a PRP network if it has been received over portA or portB of a RedBox.


For frames that originate from within the HSR ring, ~~LaneId~~ **LanId** may be set to 0 or to 1 to indicate the sending port (A = 0, B = 1).

The allocation of the 4-bit PathId field is:

- 0000 to 0001 frame from an HSR node (A/B);
- 0010 to 1111 frames from one of 7 PRP networks (A/B).

If the frame carries a VLAN tag according to IEEE 802.1Q, the HSR tag shall be inserted after the VLAN tag.

The same frame with a VLAN tag is shown in Figure 23.



IEC

Figure 23 – HSR frame with VLAN tag

The minimum frame size (including the destination MAC address and the FCS) shall be 70 octets with no VLAN tag and 74 octets with a VLAN tag. The padding shall be inserted between payload and FCS.

NOTE 1 The LSDUsize has the same definition as in PRP and includes possible padding.

NOTE 2 The reason for inserting the HSR tag after the VLAN tag is to provide faster MAC address lookup in case IVL (Independent VLAN learning as defined in IEEE 802.1Q) is used.

NOTE 3 The minimum frame size ensures that the frame will meet the ISO/IEC/IEEE 8802-3:2014 minimum frame size of 64 octets after removal of the HSR tag and of the VLAN tag.

5.7.2 HSR_Supervision frame

5.7.2.1 Sending

Each DANH shall multicast an HSR_Supervision frame over both its ports with the format specified in Table 9 or Table 10.

NOTE 1 The frame has been padded to 70 octets (no VLAN) and 74 octets (with VLAN) respectively since the LRE can remove the HSR tag (6 octets) and VLAN tag (4 octets) when passing the frame to the upper layers.

Table 9 – HSR_Supervision frame with no VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2	HSR_DestinationAddress = multicast (01-15-4E-00-01-XX)															
4									lsb							
6						msb	U/L	0								
8	SourceAddress (MAC address of the node)															
10									lsb							
12	HSR_EtherType															
14	pathId				HSR_LSDUsize = 52 (with padding)											
16	SeqNr															
18	SupEthertype															
20	SupPath				SupVersion											
22	SupSequenceNumber															
24	TLV1.Type = 23								TLV1.Length = 6							
26						msb	U/L	0								
28	MacAddress (MAC address of the DANH)															
30																
32	TLV2.Type = 30								TLV2.Length = 6							
34						msb	U/L	0								
36	RedBoxMacAddress															
38									lsb							
40	TLV0.Type = 0								TLV0.Length = 0							
64	Padding to 70 octets (64 + 6 + 0 VLAN)															
66	FCS															
68																
70																

Table 10 – HSR_Supervision frame with optional VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14			prio		cti											
16																
18																
20																
22																
24																
26																
28																
30						Msb	U/L	0								
32																
34																
36																
38						Msb	U/L	0								
40																
42									lsb							
44																
46																
68																
70																
72																
74																

NOTE 2 The format of the frame is nearly identical to that of the PRP supervision frames.

NOTE 3 The Supervision frames allow the NodesTable to keep track of the presence of other nodes identified by their MAC address. When a node does not send a Supervision frame within a given time-out (in the order of one minute), its address is removed from the NodesTables if such are used.

5.7.2.2 HSR_Supervision frame contents

The content of the HSR_Supervision frame shall be defined as follows:

HSR_DestinationAddress	Reserved multicast address 01-15-4E-00-01-XX. By default XX is "00", but if conflicts arise, XX can be configured to take any value between 0x00 and 0xFF.
SourceAddress	MAC address of the node.
VLAN tag (optional)	IEEE 802.1Q tag if used.
HSR_EtherType	Ethertype of the HSR protocol, see 5.7.1
PathId	Path over which this frame was sent, in particular, identification of the PRP network.
LSDUsize	Size of the LSDU in octets.
SeqNr	SendSeq of the sender.
SupEtherType	Ethertype for supervision frames, shall be 0x88FB
SupPath	Unused, set to 0
SupVersion	Protocol version, set to "1" (one) for this version of HSR. Implementation of version X of the protocol shall interpret version >X as if they were version X, ignoring any parameters and/or flags added by the more recent version, and interpret version <=X exactly as specified for the version concerned. The version shall not exceed the value of 64.
SupSequenceNumber	Sequence number of the supervision frames
TLV1.Type	Has a value of 23 to indicate that it is an HSR node in normal operation. Other values are reserved.
TLV1.Length	Indicates the length of the following MAC address.
MacAddress	MAC address used by the sender.
TLV2.Type	Source node identifier set to 30 if the source node is a RedBox (see 4.3.3)
TLV2.Length	Indicates the length of the following MAC address.
RedBoxMacAddress	MAC addresses used by the RedBox on its interlink.
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

5.7.2.3 Reception of a HSR_Supervision frame and NodesTable

When receiving a first HSR_Supervision frame over any ring port, a node shall create an entry in the NodesTable corresponding to the MacAddress of that source as indicated in the message body, not in the source address of the frame, and register the port over which this frame came.

Subsequent reception of frames shall be registered with the objective of identifying reception errors.

5.7.2.4 Non-Reception of a HSR_Supervision frame

When a device detects another device as non-existent, it shall, if this device resumes transmission, accept frames from that device as if that device had not been connected before.

NOTE The mechanism to detect non-existing devices is not specified, it can be assumed that reading the NodesTable resets or reads a frame counter.

5.8 Constants

The constant parameters are shown in Table 11.

NOTE Other values are allowed at the user's responsibility.

Table 11 – HSR Constants


Constant	Description	Default value
LifeCheckInterval	How often a node sends an HSR_Supervision frame	2 000 ms
NodeForgetTime	Time after which a node entry is cleared in the NodesTable	60 000 ms
ProxyNodeTableForgetTime	Time after which a node entry is cleared in the ProxyNodeTable	60 000 ms
ProxyNodeTableMaxEntries	Maximum number of entries in the ProxyNodeTable	512
EntryForgetTime	Time after which an entry is removed from the duplicates	400 ms
NodeRebootInterval	Minimum time during which a node that reboots remains silent	500 ms
MulticastFilterSize	Number of multicast addresses to be filtered	16

5.9 HSR service specification

Objects in the LRE, e.g. error counters, can be accessed from an application running on the same node using direct calls or from an application on another node using a management protocol such as SNMP or a specific protocol such as IEC 61850.

All LRE objects are described in Clause 7 as an MIB conformant to SNMP (RFC 2578, RFC 3418), which is common for PRP and HSR.

For better understanding of the MIB counters, Figure 24 shows the structure of a node with the entities that maintain the counters.



IEC

Figure 24 – HSR node with management counters

Figure 25 shows the management information base for a RedBox.




Figure 25 – HSR RedBox with management counters

6 Protocol Implementation Conformance Statement (PICS)

The PICS shall indicate if the following options are supported:

- **SNMP_MIB:** ability to support the SNMP MIB
- **NTAB_SIZ:** number of entries in the NodesTable (0 = no NodesTable)
- **PRIO_QTY:** number of supported priorities
- **VLAN_QTY:** number of supported VLANS
- **MULT_QTY:** number of filtered multicast addresses
- **CLK_1588:** support of IEC 61588 synchronization
- **PRP_SRP:** ability to perform as non-bridging node with no PRP
- **PRP_RSTP:** ability to perform as an RSTP bridge element with designated port role
- **PRP_MRP:** ability to perform as an MRP bridge element (client or master)
- **HSR_H:** ability to support HSR Mode H (default)
- **HSR_N:** ability to support HSR Mode N (no forwarding)
- **HSR_M:** ability to support HSR Mode M (mixed forwarding)
- **HSR_T:** ability to support HSR Mode T (transparent forwarding)
- **HSR_U:** ability to support HSR Mode U (forwarding also when unique destination)
- **HSR_X:** ability to support HSR Mode X (no forwarding on counter-duplicate)
- **RBX_PRP:** RedBox with PRP ports
- **RBX_HSR:** RedBox with HSR ports

- QBX_HSR: QuadBox integrating two RedBoxes
- RBX_PNT: number of entries in the ProxyNodeTable

7 PRP/HSR Management Information Base (MIB)

The MIB objects reflect the arguments of the service parameters which bear the same name, with an uppercase first letter. If the PICS option SNMP_MIB is true, the MIB data structures defined in Clause 7 shall be available at OID = 1.0.62439 in addition to the MIBs that the ports may provide, with the following definition.

NOTE When copying the MIB from a PDF document, all headers and footers are removed and characters are converted to the character format that SNMP requires. Although care has been taken not to use illegal characters, some can appear in the final document due to the editing process.

```
-- *****
IEC-62439-3-MIB DEFINITIONS ::= BEGIN
-- *****
-- Imports
-- *****
IMPORTS

    MODULE-IDENTITY, OBJECT-TYPE,
    Counter32, TimeTicks, Integer32, Unsigned32          FROM SNMPv2-SMI
    OBJECT-GROUP,
    MODULE-COMPLIANCE                                   FROM SNMPv2-CONF
    TruthValue, RowStatus, MacAddress, DisplayString,
    TEXTUAL-CONVENTION                                  FROM SNMPv2-TC;

-- *****
-- Root OID
-- *****

iec62439 MODULE-IDENTITY

    LAST-UPDATED "201202170000Z" -- February 17, 2012 "201405220000Z" -- 2014, May 22
    ORGANIZATION "IEC/SC 65C"
    CONTACT-INFO "
        International Electrotechnical Commission
        IEC Central Office
        3, rue de Varembe
        P.O. Box 131
        CH-1211 GENEVA 20
        Switzerland
        Phone: +41 22 919 02 11
        Fax: +41 22 919 03 00
        email: info@iec.ch
    "

    DESCRIPTION "
        This MIB module defines the Network Management interfaces
        for the redundancy protocols defined by the IEC 62439 suite.
        This MIB exposes the IEC62439-3 objects (PRP + HSR)
    "

    REVISION "201405220000Z" -- 2014, May 22
    DESCRIPTION "
        added reference to ptp MIB and lreDupListResideMaxTime and type
        SecondFraction
    "

    REVISION "201202170000Z" -- February 17, 2012
    DESCRIPTION "
        Consistency brought into line to mrp, crp, brp MIBs
    "

    REVISION "201108260000Z" -- August 26, 2011
    DESCRIPTION "
```

This MIB is aligned with the changes to PRP and HSR as defined in the Amendment to IEC 62439-3

REVISION "200811100000Z" -- November 10, 2008
DESCRIPTION "

Separation of IEC 62439 into a suite of documents.
This MIB applies to IEC 62439-3, added HSR functionality
"

REVISION "200612160000Z" -- December 16, 2006
DESCRIPTION "

Initial version of the Network Management interface for the
Parallel Redundancy Protocol
"

::= {iso std(0) 62439 }

-- *****
-- Redundancy Protocols
-- *****

mrp OBJECT IDENTIFIER ::= { iec62439 1 }
prp OBJECT IDENTIFIER ::= { iec62439 2 }
crp OBJECT IDENTIFIER ::= { iec62439 3 }
brp OBJECT IDENTIFIER ::= { iec62439 4 }
drp OBJECT IDENTIFIER ::= { iec62439 5 }
rrp OBJECT IDENTIFIER ::= { iec62439 6 }
ptp OBJECT IDENTIFIER ::= { iec62439 7 }

-- *****
-- Textual conventions
-- *****

SecondFraction ::= TEXTUAL-CONVENTION

DISPLAY-HINT "d"

STATUS current

DESCRIPTION

"time interval expressed in multiple of 2**⁻¹⁶ = 15,7 microseconds
this corresponds to the fraction of seconds in the NTP time representation
the minimum time is 0 microseconds, the maximum 18 hours
"

REFERENCE "IEC 62439-3"

SYNTAX Integer32

-- *****
-- Objects of the PRP Network Management
-- *****

linkRedundancyEntityNotifications OBJECT IDENTIFIER ::= { prp 20 }
linkRedundancyEntityObjects OBJECT IDENTIFIER ::= { prp 21 }
linkRedundancyEntityConformance OBJECT IDENTIFIER ::= { prp 22 }

-- *****
lreConfiguration OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 0 }
lreStatistics OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 1 }
lreConfigurationGeneralGroup OBJECT IDENTIFIER ::= { lreConfiguration 0 }
lreConfigurationInterfaceGroup OBJECT IDENTIFIER ::= { lreConfiguration 1 }
lreStatisticsInterfaceGroup OBJECT IDENTIFIER ::= { lreStatistics 1 }

-- *****
-- Objects for lreConfigurationGeneralGroup
-- *****

lreManufacturerName OBJECT-TYPE
SYNTAX DisplayString
MAX-ACCESS read-only
STATUS current
DESCRIPTION "specifies the name of the LRE device manufacturer"
::= { lreConfigurationGeneralGroup 1 }


```

lreInterfaceCount OBJECT-TYPE
    SYNTAX Integer32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "total number of LREs present in this system."
 ::= { lreConfigurationGeneralGroup 2 }

-- *****
-- Objects for lreConfigurationInterfacesGroup
-- *****

lreConfigurationInterfaces OBJECT IDENTIFIER
 ::= { lreConfigurationInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesConfigTable***
-- *****

lreInterfaceConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        list of PRP/HSR LREs. Each entry corresponds
        to one PRP/HSR Link Redundancy Entity (LRE), each representing a
        pair of LAN ports A and B. Basic devices supporting PRP/HSR may
        have only one LRE and thus one entry in the table, while more complex
        devices may have several entries for multiple LREs.
        "
 ::= { lreConfigurationInterfaces 1 }

lreInterfaceConfigEntry OBJECT-TYPE
    SYNTAX LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        each entry contains management information applicable to a
        particular LRE.
        "
    INDEX { lreInterfaceConfigIndex }
 ::= { lreInterfaceConfigTable 1 }

LREInterfaceConfigEntry ::=
    SEQUENCE {
        lreInterfaceConfigIndex Unsigned32,
        lreRowStatus RowStatus,
        lreNodeType INTEGER,
        lreNodeName DisplayString,
        lreVersionName OCTET STRING,
        lreMacAddress MacAddress,
        lrePortAdminStateA INTEGER,
        lrePortAdminStateB INTEGER,
        lreLinkStatusA INTEGER,
        lreLinkStatusB INTEGER,
        lreDuplicateDiscard INTEGER,
        lreTransparentReception INTEGER,
        lreHsrLREMode INTEGER,
        lreSwitchingEndNode INTEGER,
        lreRedBoxIdentity INTEGER,
        lreEvaluateSupervision TruthValue,
        lreNodesTableClear INTEGER,
        lreProxyNodeTableClear INTEGER,
        lreDupListResideMaxTime SecondFraction
    }

lreInterfaceConfigIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible

```

```

STATUS current
DESCRIPTION "A unique value for each LRE."
::= { lreInterfaceConfigEntry 1 }

lreRowStatus OBJECT-TYPE
SYNTAX RowStatus
MAX-ACCESS read-create
STATUS current
DESCRIPTION "indicates the status of the LRE table entry"
::= { lreInterfaceConfigEntry 2 }

lreNodeType OBJECT-TYPE
SYNTAX INTEGER
{
    prpmodel (1),
    hsr (2)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION "
    specifies the operation mode of the LRE:

    PRP mode 1 (1)
    HSR mode (2)

    Note: PRP mode 0 is considered deprecated and is not supported by this
           revision of the MIB
    "
::= { lreInterfaceConfigEntry 3 }

lreNodeName OBJECT-TYPE
SYNTAX DisplayString
MAX-ACCESS read-write
STATUS current
DESCRIPTION "specifies this LRE's node name"
::= { lreInterfaceConfigEntry 4 }

lreVersionName OBJECT-TYPE
SYNTAX OCTET STRING (SIZE(1..32))
MAX-ACCESS read-only
STATUS current
DESCRIPTION "specifies the version of this LRE's software"
::= { lreInterfaceConfigEntry 5 }

lreMacAddress OBJECT-TYPE
SYNTAX MacAddress
MAX-ACCESS read-write
STATUS current
DESCRIPTION "
    Specifies the MAC address to be used by this LRE. MAC
addresses are
    identical for all ports of a single LRE
    "
::= { lreInterfaceConfigEntry 6 }

lrePortAdminStateA OBJECT-TYPE
SYNTAX INTEGER
{
    notActive (1),
    active (2)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION "
    Specifies whether the port A shall be active or not Active
through

```

```

        administrative action (Default: active).
        "
::= { lreInterfaceConfigEntry 7 }

lrePortAdminStateB OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        notActive (1),
        active (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Specifies whether the port B shall be active or not Active
through
        administrative action (Default: active).
        "
::= { lreInterfaceConfigEntry 8 }

lreLinkStatusA OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port A"
::= { lreInterfaceConfigEntry 9 }

lreLinkStatusB OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port B"
::= { lreInterfaceConfigEntry 10 }

lreDuplicateDiscard OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        doNotDiscard (1),
        discard (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        specifies whether a duplicate discard algorithm is used at
reception
        (Default: discard).
        "
::= { lreInterfaceConfigEntry 11 }

lreTransparentReception OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        removeRCT (1),
        passRCT (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "

```

```

        if removeRCT is configured, the RCT is removed when
forwarding to the upper layers, only applicable for PRP LRE (Default: removeRCT).
    "
::= { lreInterfaceConfigEntry 12 }

```

lreHsrLREMode OBJECT-TYPE

```

SYNTAX          INTEGER
{

```

```

    modeh (1),
    moden (2),
    modet (3),
    modeu (4),
    modem (5)
}

```

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"

This enumeration is only applicable if the LRE is an HSR bridging node or RedBox.

It shows the mode of the HSR LRE:

(1) Default mode: The HSR LRE is in mode h and bridges tagged HSR traffic

(2) Optional mode: The HSR LRE is in mode n and bridging between its HSR ports is disabled.

Traffic is HSR tagged.

(3) Optional mode: The HSR LRE is in mode t and bridges non-tagged HSR traffic between its HSR ports

(4) Optional mode: The HSR LRE is in mode u and behaves like in mode h, except it does not remove unicast messages

(5) Optional mode: The HSR LRE is configured in mixed mode. HSR frames are handled according to mode h. Non-HSR frames are handled according to

802.1D bridging rules.

"

```

::= { lreInterfaceConfigEntry 13}

```

lreSwitchingEndNode OBJECT-TYPE

```

SYNTAX          INTEGER
{

```

```

    nonbridgingnode(1),
    bridgingunspecified(2),
    prpnode(3),
    hsrredboxsan(4),
    hsrnode(5),
    hsrredboxhsr(6),
    hsrredboxprpa(7),
    hsrredboxprpb(8)
}

```

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"

This enumeration shows which feature is enabled in this particular LRE:

(1): an unspecified non-bridging node, e.g. SRP.

(2): an unspecified bridging node, e.g. RSTP.

(3): a PRP node/RedBox.

(4): an HSR RedBox with regular Ethernet traffic on its interlink.

(5): an HSR switching node.

(6): an HSR RedBox with HSR tagged traffic on its interlink.

(7): an HSR RedBox with PRP traffic for LAN A on its interlink.

(8): an HSR RedBox with PRP traffic for LAN B on its interlink.

"

```

::= { lreInterfaceConfigEntry 14 }

```

lreRedBoxIdentity OBJECT-TYPE

```

SYNTAX          INTEGER
{

```

```

    id1a (2),
    id1b (3),

```

```

id2a (4),
id2b (5),
id3a (6),
id3b (7),
id4a (8),
id4b (9),
id5a (10),
id5b (11),
id6a (12),
id6b (13),
id7a (14),
id7b (15)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION
    "
    Applicable to RedBox HSR-PRP A and RedBox HSR-PRP B. One ID
    is used by
    one pair of RedBoxes (one configured to A and one configured to B)
    coupling an HSR ring to a PRP network. The integer value states the value
    of the path field a RedBox inserts into each frame it receives from its
    interlink and injects into the HSR ring. When interpreted as binary values,
    the LSB denotes the configuration of the RedBox (A or B), and the following
    3 bits denote the identifier of a RedBox pair.
    "
::= { lreInterfaceConfigEntry 15}

lreEvaluateSupervision OBJECT-TYPE
    SYNTAX TruthValue
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        True if the LRE evaluates received supervision frames. False
        if it
        drops the supervision frames without evaluating. Note: LREs are required
        to send supervision frames, but reception is optional. Default value is dependent
        on implementation.
        "
::= { lreInterfaceConfigEntry 16}

lreNodesTableClear OBJECT-TYPE
    SYNTAX INTEGER
    {
        noOp (0),
        clearNodeTable (1)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "specifies that the Node Table is to be cleared"
::= { lreInterfaceConfigEntry 17}

lreProxyNodeTableClear OBJECT-TYPE
    SYNTAX INTEGER
    {
        noOp (0),
        clearProxyNodeTable (1)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "specifies that the Proxy Node Table is to be cleared"
::= { lreInterfaceConfigEntry 18}

lreDupListResideMaxTime OBJECT-TYPE
    SYNTAX SecondFraction
    UNITS "binaryFractionOfSecond"
    MAX-ACCESS read-write
    STATUS current

```

```

    DESCRIPTION          "the longest time an entry may reside in the duplicates list,
expressed as the number of seconds multiplied by 65536;
                           the default value is 26214 x 15 us, or 400 ms; too low a value can
cause broadcast storms"
    DEFVAL                {26214}
    ::= { lreInterfaceConfigEntry 19}

-- *****
-- *** End lreInterfaceConfigTable ***
-- *****

-- *****
-- Objects for lreStatisticsInterfacesGroup
-- *****

lreStatisticsInterfaces OBJECT IDENTIFIER
::= { lreStatisticsInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesStatsTable ***
-- *****

lreInterfaceStatsTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LREInterfaceStatsEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        list of PRP/HSR LREs. Each entry corresponds
        to one PRP/HSR Link Redundancy Entity (LRE), each representing a
        pair of LAN ports A and B and a port C towards the application/interlink.
        Basic devices supporting PRP/HSR may have only one LRE and thus one entry in
        the table, while more complex devices may have several
        entries for multiple LREs.
        "
    ::= { lreStatisticsInterfaces 1 }

lreInterfaceStatsEntry OBJECT-TYPE
    SYNTAX LREInterfaceStatsEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        An entry containing management information applicable to a
        particular LRE.
        "
    INDEX { lreInterfaceStatsIndex }
    ::= { lreInterfaceStatsTable 1 }

LREInterfaceStatsEntry ::=
    SEQUENCE {
        lreInterfaceStatsIndex Unsigned32,
        lreCntTxA Counter32,
        lreCntTxB Counter32,
        lreCntTxC Counter32,
        lreCntErrWrongLanA Counter32,
        lreCntErrWrongLanB Counter32,
        lreCntErrWrongLanC Counter32,
        lreCntRxA Counter32,
        lreCntRxB Counter32,
        lreCntRxC Counter32,
        lreCntErrorsA Counter32,
        lreCntErrorsB Counter32,
        lreCntErrorsC Counter32,
        lreCntNodes Integer32,
        lreCntProxyNodes Integer32,
        lreCntUniqueA Counter32,
        lreCntUniqueB Counter32,
    }

```

```

    lreCntUniqueC Counter32,
    lreCntDuplicateA Counter32,
    lreCntDuplicateB Counter32,
    lreCntDuplicateC Counter32,
    lreCntMultiA Counter32,
    lreCntMultiB Counter32,
    lreCntMultiC Counter32,
    lreCntOwnRxA Counter32,
    lreCntOwnRxB Counter32
}

lreInterfaceStatsIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each LRE."
 ::= { lreInterfaceStatsEntry 1 }

lreCntTxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port A that are HSR tagged or
        fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 2 }

lreCntTxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port B that are HSR tagged or
        fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 3 }

lreCntTxC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent towards the application interface of
        the DANP or DANH
        or over the interlink of the RedBox. Frames with and without PRP RCT or HSR
        tag are counted, but not link-local frames.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 4 }

lreCntErrWrongLanA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "

```

```

        number of frames with the wrong LAN identifier received on
LRE port A.
        Initial value = 0. Only applicable to PRP ports.
        "
::= { lreInterfaceStatsEntry 5 }

lreCntErrWrongLanB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames with the wrong LAN identifier received on
LRE port B
        Initial value = 0. Only applicable to PRP ports.
        "
::= { lreInterfaceStatsEntry 6 }

lreCntErrWrongLanC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames with the wrong LAN identifier received on
the interlink of
        a RedBox. Only applicable to HSR RedBoxes in HSR-PRP configuration
        (hsrredboxprpa and hsrredboxprpb).
        "
::= { lreInterfaceStatsEntry 7 }

lreCntRxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames received on a LRE port A. Only frames that
are HSR tagged
        or fitted with a PRP Redundancy Control Trailer are counted. Frames that are
        not forwarded anywhere (e.g. because the sender of the frame is in the proxy
        node table) are counted, too. Only frames received completely and without
        error are counted.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 8 }

lreCntRxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames received on a LRE port B. Only frames that
are HSR tagged
        or fitted with a PRP Redundancy Control Trailer are counted. Frames that are
        not forwarded anywhere (e.g. because the sender of the frame is in the proxy
        node table) are counted, too. Only frames received completely and without
        error are counted.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 9 }

lreCntRxC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "

```


number of frames received from the application interface of a DANP or DANH or

the number of number of frames received on the interlink of a RedBox. Frames with and without PRP RCT or HSR tag are counted, but not link-local frames. Only frames received completely and without error are counted.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 10 }

lreCntErrorsA OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames with errors received on this LRE port A.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 11 }

lreCntErrorsB OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames with errors received on this LRE port B.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 12 }

lreCntErrorsC OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames with errors received on the application

interface of a DANP or

DANH or on the interlink of a RedBox.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 13 }

lreCntNodes OBJECT-TYPE

SYNTAX Integer32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"number of nodes in the Nodes Table."

::= { lreInterfaceStatsEntry 14 }

lreCntProxyNodes OBJECT-TYPE

SYNTAX Integer32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of nodes in the Proxy Node Table. Only applicable to

RedBox.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 15 }

lreCntUniqueA OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

```

        number of entries in the duplicate detection mechanism on
port A for which
    no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 16 }

lreCntUniqueB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 17 }

lreCntUniqueC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 18 }

lreCntDuplicateA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port A for which
    one single duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 19 }

lreCntDuplicateB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    one single duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 20 }

lreCntDuplicateC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which one single duplicate was
received.
        Initial value = 0.

```

```

"
::= { lreInterfaceStatsEntry 21 }

lreCntMultiA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port A for which
    more than one duplicate was received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 22 }

lreCntMultiB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    more than one duplicate was received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 23 }

lreCntMultiC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which more than one duplicate was
received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 24 }

lreCntOwnRxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of HSR tagged frames received on Port A that
originated from this
device. Frames originate from this device if the source MAC matches the
MAC of the LRE, or if the source MAC appears in the proxy node table (if
implemented). Applicable only to HSR.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 25 }

lreCntOwnRxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of HSR tagged frames received on Port B that
originated from this
device. Frames originate from this device if the source MAC matches the
MAC of the LRE, or if the source MAC appears in the proxy node table (if
implemented). Applicable only to HSR.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 26 }

```

```
-- *****
-- *** End LRE InterfacesStatsTable ***
-- *****

-- *****
-- ***Begin LRE NodesTable ***
-- *****

lreNodesTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LRENodesEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        The node table (if it exists on that node) contains
information about
        all remote LRE, which advertised themselves through
supervision frames
        "
    ::= { lreStatisticsInterfaces 2 }

lreNodesEntry OBJECT-TYPE
    SYNTAX LRENodesEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        Each entry in the node table (if it exists) contains
information about
        a particular remote LRE registered in the node table, which advertised itself
        through supervision frames.
        "
    INDEX { lreInterfaceStatsIndex, lreNodesIndex }
    ::= { lreNodesTable 1 }

LRENodesEntry ::=
    SEQUENCE {
        lreNodesIndex Unsigned32,
        lreNodesMacAddress MacAddress,
        lreTimeLastSeenA TimeTicks,
        lreTimeLastSeenB TimeTicks,
        lreRemNodeType INTEGER
    }

lreNodesIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "Unique value for each node in the LRE's node table."
    ::= { lreNodesEntry 1 }

lreNodesMacAddress OBJECT-TYPE
    SYNTAX MacAddress
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "Each MAC address corresponds to a single Dual Doubly Attached Node"
    ::= { lreNodesEntry 2 }

lreTimeLastSeenA OBJECT-TYPE
    SYNTAX TimeTicks
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        Time in TimeTicks (1/100s) since the last frame from this
remote LRE was
        received over LAN A. Initialized with a value of 0 upon node registration
        in the node table.
        "

```

```
::= { lreNodesEntry 3 }
```

```
lreTimeLastSeenB OBJECT-TYPE
```

```
    SYNTAX TimeTicks
```

```
    MAX-ACCESS read-only
```

```
    STATUS current
```

```
    DESCRIPTION
```

```
        "Time in TimeTicks (1/100s) since the last frame from this
```

```
remote LRE was
```

```
    received over LAN B. Initialized with a value of 0 upon node registration
    in the node table.
```

```
"
```

```
::= { lreNodesEntry 4 }
```

```
lreRemNodeType OBJECT-TYPE
```

```
    SYNTAX INTEGER
```

```
    {
```

```
        danp (0),
```

```
        redboxp (1),
```

```
        vdanp (2),
```

```
        danh (3),
```

```
        redboxh (4),
```

```
        vdanh (5)
```

```
    }
```

```
    MAX-ACCESS read-only
```

```
    STATUS current
```

```
    DESCRIPTION
```

```
        "DAN type, as indicated in the received supervision frame"
```

```
::= { lreNodesEntry 5 }
```

```
-- *****
```

```
-- *** End LRE NodesTable ***
```

```
-- *****
```

```
-- *****
```

```
-- *** Begin LRE ProxyNodeTable ***
```

```
-- *****
```

```
lreProxyNodeTable OBJECT-TYPE
```

```
    SYNTAX SEQUENCE OF LREProxyNodeEntry
```

```
    MAX-ACCESS not-accessible
```

```
    STATUS current
```

```
    DESCRIPTION
```

```
        "The proxy node table (if implemented) contains information
```

```
about all
```

```
nodes, for which the LRE acts as a connection to the HSR/PRP
```

```
network.
```

```
"
```

```
::= { lreStatisticsInterfaces 3 }
```

```
lreProxyNodeEntry OBJECT-TYPE
```

```
    SYNTAX LREProxyNodeEntry
```

```
    MAX-ACCESS not-accessible
```

```
    STATUS current
```

```
    DESCRIPTION
```

```
        "Each entry in the proxy node table contains information about
```

```
a particular nodefor
```

```
which the LRE acts as a connection to the HSR/PRP network.
```

```
"
```

```
    INDEX { lreInterfaceStatsIndex, lreProxyNodeIndex }
```

```
::= { lreProxyNodeTable 1 }
```

```
LREProxyNodeEntry ::=
```

```
    SEQUENCE {
```

```
        lreProxyNodeIndex Unsigned32,
```

```
        lreProxyNodeMacAddress MacAddress
```

```
    }
```

```

lreProxyNodeIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each node in the LRE's proxy node table."
    ::= { lreProxyNodeEntry 1 }

lreProxyNodeMacAddress OBJECT-TYPE
    SYNTAX MacAddress
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        Each entry contains information about a particular node
        for which the LRE acts as a proxy for the HSR/PRP network.
        "
    ::= { lreProxyNodeEntry 2 }

-- *****
-- *** End LRE ProxyNodeTable ***
-- *****

*****

=====
-- Conformance Information
=====

linkRedundancyConformance OBJECT IDENTIFIER ::= { linkRedundancyEntityConformance 1 }
lreGroups OBJECT IDENTIFIER ::= { linkRedundancyConformance 1 }
lreDefaultGrp OBJECT-GROUP
    OBJECTS {
        lreManufacturerName,
        lreInterfaceCount,
        lreRowStatus,
        lreNodeType,
        lreNodeName,
        lreVersionName,
        lreMacAddress,
        lrePortAdminStateA,
        lrePortAdminStateB,
        lreLinkStatusA,
        lreLinkStatusB,
        lreDuplicateDiscard,
        lreTransparentReception,
        lreHsrLREMode,
        lreSwitchingEndNode,
        lreRedBoxIdentity,
        lreEvaluateSupervision,
        lreNodesTableClear,
        lreProxyNodeTableClear,
        lreDupListResideMaxTime,
        lreCntTxA,
        lreCntTxB,
        lreCntTxC,
        lreCntErrWrongLanA,
        lreCntErrWrongLanB,
        lreCntErrWrongLanC,
        lreCntRxA,
        lreCntRxB,
        lreCntRxC,
        lreCntErrorsA,
        lreCntErrorsB,
        lreCntErrorsC,
        lreCntNodes,
        lreCntProxyNodes,
        lreCntUniqueA,
        lreCntUniqueB,
        lreCntUniqueC,
        lreCntDuplicateA,
        lreCntDuplicateB,
        lreCntDuplicateC,
    }

```

```

    lreCntMultiA,
    lreCntMultiB,
    lreCntMultiC,
    lreCntOwnRxA,
    lreCntOwnRxB,
    lreNodesMacAddress,
    lreTimeLastSeenA,
    lreTimeLastSeenB,
    lreRemNodeType,
    lreProxyNodeMacAddress
  }
  STATUS current
  DESCRIPTION
    "Objects in the default group"
  ::= { lreGroups 1 }

-- *****
=====
-- MIB module Compliance statements
=====

linkRedundancyCompliances OBJECT IDENTIFIER ::= { linkRedundancyEntityConformance 2 }
linkRedundancyCompliance MODULE-COMPLIANCE
  STATUS current
  DESCRIPTION "Compliance for support by IEC 62439-3 module"
  MODULE
  MANDATORY-GROUPS {

    lreDefaultGrp
  }
  ::= { linkRedundancyCompliances 1 }

END

```

Annex A (normative)

Use of IEC 61588 and IEEE C37.238 for IEC 62439-3

A.1 Overview

This Annex specifies the mapping of the Precision Time Protocol (IEC 61588:2009) time distribution to PRP and HSR, using specifically the IEEE C37.238 Standard Profile for Use of IEEE 1588™ Precision Time Protocol in Power System Applications.

This Annex has been transferred from the IEC 61850-90-4 Technical Report (Network Engineering Guidelines for Substation Automation) which will now refer to this document.

A.2 PTP profile for Automation Networks using PRP & HSR

IEC 61588 (PTP) allows numerous options. Only a subset is needed to cover automation networks using HSR and PRP.

The following profile of IEC 61588 aims at reaching a precision of 2,5 µs in a maximum size network consisting of up to 50 cascaded Transparent Clocks, each using a local oscillator with an accuracy better than 50 ppm and a drift of less than 1 ppm/s.

- All network elements (nodes, bridges) are IEC 61588-aware.
(all bridges between clocks can correct the clock, routers are not considered)
- Media converters (e.g. fibre – copper) introduce less than 50 ns of jitter each.
- Cable asymmetry over one link does not exceed 25 ns.
- Each Transparent Clock introduces less than 50 ns of clock inaccuracy in the correction field of forwarded Sync messages.
- Boundary Clocks are not considered.
(Boundary Clocks are located at the outskirts of the multicast domain).
- Only Ethernet layer 2 communication applies.
(UDP communication is not used)
- Only multicast PTP messages are used.
(unicast messages are not used)
- Only peer-to-peer delays are calculated.
(end-to-end path delay computation is not used)
- Both 1 step and 2 step synchronization are allowed.
- Best Master Clock contest applies to all clocks except slave-only;
- Default settings are specified as:
 - Sync messages period (logAnnounceInterval): once every 1 s;
 - Announce message period (logSyncInterval): once every 1 s;
 - Timeout for Announce messages (announceReceiptTimeout): 2 s, resp. 3 s;
 - Pdelay_Req period (logMinDelayReqInterval): once every 1 s;
 - Holdover time: accuracy kept within limits after signal loss of at most 5s.
 - Priority1 is 128 for Grandmaster capable devices, 255 for slave-only devices
 - Priority2 is 128 for Grandmaster capable devices, 255 for slave-only devices
 - The Grandmaster clockClass is 6 (normal) or 7 (degraded)

A.3 PTP operation (informative)

A.3.1 Node types

PTP distinguishes the following node types:


- Grand Master Clock (GMC), which is the clock connected to a reference signal (GPS or atomic clock) and that multicasts the time through Sync messages.
- Master Clocks (MC) that multicast the time within a subdomain, but are not connected to a reference clock.
- Master-capable clocks that can become master or grand-master.
- Transparent Clocks (TC) that forward the Sync messages and correct the time by evaluating the peer delay and residence delay of Sync messages.
- Ordinary Clocks (OC) that can be either slave or master clocks.
- Hybrid Clocks (HC) that combine a Transparent Clock and an Ordinary Clock.
- Boundary Clocks (BC) that have both a slave port (towards the GMC) and a master clock (towards OCs). A boundary clock is synchronized by a master clock and acts as a master for an own clock island. Within this island, it does not forward the messages received by the GMC, but sends its own PTP messages, possibly with a different period.

A.3.2 Principle and one-step operation

The primary time master, called GrandMaster Clock (GMC) periodically (e.g. every 1 s) multicasts a Sync message containing the reference time. This Sync message transits through the network over links and TCs. The Sync message suffers network delays consisting of the link delays λ between nodes and of the residence delays ρ in the nodes, before it reaches the slave clocks. The TCs therefore compute a correction for the link and residence delays.

Figure A.1 shows the one-step PTP clock distribution to explain the principle.

To calculate the link delays, each node first sends peer delay request messages (Pdelay_Req) to all its neighbours, even in the direction opposite to the current GMC since a redundant GMC may exist. The neighbour responds with a Pdelay_Resp message and the link delay is calculated as:



IEC 1153/12

Figure A.1 – PTP one-step clock synchronization and delay measurement

In one-step clocks, each TC j adds its correction ($\lambda_{kj} + \rho_j$ + received correction from clock k) to the Sync message's correction field on-the-fly, while the frame is being transmitted on the egress port. This method requires a dedicated hardware to modify the correction field and the checksum of the outgoing Sync frame while parts of the message are already been transmitted (see Figure A.1).

A.3.3 Two-step operation

To avoid a dedicated hardware for modifying the Sync message on-the-fly, a two-steps clock records the precise time at which the Sync was received, forwards the Sync message without modification (but announcing a Follow_Up message) and records the precise time at which the Sync was sent. There exist dedicated PHY chips that time-stamp messages within a few ns, so sending the Sync message is not time-critical, it can be done by a CPU.

The two-step TCs calculates the correction field as the sum of the ingress Sync correction, the residence delay and the (previously measured) peer delay and sends a Follow_Up message as Figure A.2 shows. The same scheme applies to the calculation of the peer delay.

The receiver assumes that Sync and Follow_Up messages are paired; they have the same PTP sequence number.

NOTE PRP and HSR present a special case, since Sync, Follow_Up and Announce are paired on each port separately.




Figure A.2 – PTP two-step clock synchronization and delay measurement

The engineer is faced with two synchronization methods, depending on the ability of its devices. Commercial clocks use mainly two-step PTP, few offer one-step clocks. The two systems can be mixed if devices accept both one-step and two-step clocks. Usually, two-steps are preferred in the RSTP and PRP networks, while one-step is preferred in HSR rings.

A.3.4 Two-step to one-step translation

Both one-step TCs and two-step TCs are allowed in arbitrary mixture in PTP networks.

For HSR, one-step clock is preferred. A RedBox converts two-step PTP messages coming over a PRP path to one-step PTP messages forwarded to the HSR network.

Figure A.3 shows the conversion from two-step to one-step in a chain of one-step and two-step transparent clocks between an ordinary clock (slave) and the grandmaster clock (master).




Figure A.3 – Two-step and one-step transparent clocks translator

A sender is either a one step or a two step clock; a receiver is able to deal with both one step and two step. A one step transparent clock will use a translator to convert from two step to one step.

A node translates the two step to one step by awaiting the Follow_Up before forwarding the one step Sync, which then carries the sum of the correction fields of the two step Sync, of the Follow_Up, the residence delay of node and the estimated delay to the next peer, as Figure A.4 shows.




Figure A.4 – Two-step to one-step translation

On the reverse path, no such translation is needed since a node can send one step Syncs to the next LAN device, which will then translate it to a two step clock if needed.

A.4 — PRP Mapping

A.4.1 — Operation in PRP

The operating mode of PRP is that a DANP is attached to two distinct LANs. A DANP source sends over both LANs the same frame after appending the RCT. A frame is not modified while

transiting through the LANs. The DANP destination receives the first frame of a pair and discards the duplicate on the base of its RCT and source address.


This operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same in the two LANs and some frames are modified in the bridges while transiting through the LAN. A DANP does not receive the same PTP message from both ports.

Figure A.5 shows the connection between a GMC and an OC in a PRP network, both as DANP. The path of a Sync message through the LANs with the correction of the clock for the residence and peer-to-peer delays is suggested using a 2-step clock as an example.

When applied to PTP, the PRP frames are treated differently:

- Pdelay_Req, Pdelay_Resp and Pdelay_Resp_Follow_Up carry no RCT, since these messages are link-specific.
- One-step Syncs frames are modified by TCs to adjust the correction field; implementations are not obliged to forward the part coming after the payload.
- Follow_Up frames are regenerated by every two-step clock and carry no RCT.
- BCs present in the LAN would generate their own Announce and Sync frames with no RCT appended.
- Announce frames can also be modified by TCs if the TLVs contain dynamic information, such as IEEE C37.238 foresees for experimental purpose.

Therefore, the RCT cannot be relied upon for PTP messages.



IEC 1157/12

Figure A.5 – Connection of a Grandmaster Clock to an Ordinary Clock over PRP

~~A.4.2 — Transmission through the network~~

~~One step or two step clocks may be implemented with PRP.~~

~~A.4.3 — GMC as DANP~~

~~When operating on a PRP network, a GMC shall be either:~~

- ~~• Connected through a RedBox or~~
- ~~• Attached directly as a DANP.~~

~~The GMC shall send Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up and Follow_Up without an RCT.~~

~~NOTE A PRP topology in which one GMC is located on LAN A and another on LAN B should be avoided, since a GMC on one LAN ignores another GMC on the other LAN and cannot serve as a back-up.~~

~~A.4.4 — Reception in the DANPs~~

~~A DANP shall implement an Ordinary Clock.~~

~~A DANP shall send Pdelay_Req, Pdelay_Resp, and Pdelay_Resp_Follow_Up without an RCT.~~

~~A DANP shall not rely on their RCT for received PTP messages.~~

~~A DANP shall consider both ports independently, as if there were independent clock sources.~~

~~A DANP shall handle all PTP messages for each port separately.~~

~~A DANP receives from each port two Announce messages coming from the same grandmaster with different delays. When it receives an Announce message, a port shall await the Announce message on the other port (or proceed if Announce does not arrive within the Announce time-out) before taking a decision from which port to take its synchronization.~~

~~A port shall apply the Best Master Clock algorithm on port A and port B.~~

~~If the Announce messages from both ports have the same Clock Identity (they come from the same GMC), a DANP shall use Syncs from both ports for synchronization, else a DANP shall listen to the port that provides the best clock quality.~~

~~A DANP shall accept both one step and two step messages. If it receives a two step Sync, a port shall await the Follow_Up message from the same port before using the Sync timestamps to adjust its clock. If it receives a two step Pdelay_Resp, a port shall await the Pdelay_Resp_Follow_Up from that same direction to calculate the path delay over that port.~~

~~A.4.5 — RedBox connection~~

~~A PRP redundancy box (RedBox) shall implement a peer-to-peer Transparent Clock and may include an Ordinary Clock if the local applications require synchronization.~~


~~A Boundary Clock in a RedBox is allowed as a means to connect external devices outside of the PRP — HSR domain.~~

A.5 HSR Mapping

A.5.1 PTP traffic in HSR

The operating mode of HSR is that a DANH sends the same frame in both directions into the ring, prefixed with the HSR Tag containing a sequence number. Each node relays the frame until it reaches the source again or is otherwise discarded. A destination node therefore receives two frames, uses the first and discards the duplicate.

This operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same going clockwise or counter-clockwise in the ring. A DANH does not receive the same message from both ports since some messages are modified by intermediate nodes. Figure A.6 illustrates the paths of the different PTP messages in an HSR ring network.



IEC 1158/12

Figure A.6 – HSR with one GMC

For one-step operation, each DANH implements the one-step Sync and peer-to-peer delay mechanism and exchanges the messages shown in Figure A.7.




Figure A.7 – PTP messages sent and received by an HSR node (one-step)

Alternatively, the HSR nodes may operate as two-step clocks as Figure A.8 shows.




Figure A.8 – PTP messages sent and received by an HSR node (two-step)

A.5.2 HSR nodes

A DANH acting as grandmaster clock shall send the Announce and Sync messages in both directions with an HSR tag.

The Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up and Follow_Up messages shall not contain an HSR tag.

A DANH shall implement a Hybrid Clock, made of a Transparent Clock and an optional Ordinary Clock.

A DANH shall not introduce a delay of more than 125 µs and a time inaccuracy of more than 50 ns when forwarding a Sync message.

A DANH shall consider both ports independently, as if there were independent clock sources.

~~A DANH shall handle all PTP messages for each port separately.~~

~~A DANH receives from each port two Announce messages coming from the same grandmaster with the same Clock Identity and different delays.~~

~~If it receives an Announce message on one port, a node shall await the Announce message on the other port (or proceed if Announce does not arrive within the Announce time-out) before taking a decision from which port to take its synchronization.~~

~~A port shall apply the Best Master Clock algorithm on port A and port B.~~

~~If the Announce messages from both ports have the same Clock Identity, a DANP shall use Syncs from both ports for synchronization, else a DANH shall listen to the port that provides the best clock quality.~~

~~A DANH receives from each port two Sync messages coming from the same grandmaster with different correction fields.~~

~~When it receives a two-step Sync, a port shall await the Follow-Up message from the same port before using the Sync timestamps to adjust its clock.~~

~~If it receives a two-step Pdelay_Resp, a port shall await the Pdelay_Resp_Follow-Up from that same direction to calculate the path delay over that port.~~

~~NOTE Both Sync can come apart with a delay of up to 5 ms in a large ring.~~

~~A.5.3 Redundant clocks in HSR~~

~~As Figure A.6 shows, there can be a redundant master clock in the ring. However, both cannot be active at the same time in RTP. Therefore, the grandmaster clock is selected among the master-capable clocks according to the Best Master Clock of PTP.~~

~~In case of loss of the master, the back-up master detects the loss and start broadcasting its own Announce messages.~~

~~A.5.4 Attachment of the GMC to an external LAN by a Transparent Clock~~

~~A GMC may be connected to an HSR ring through a RedBox (Figure A.9).~~

~~In its simplest form, this RedBox can be a three-port HSR node connected through an interlink port to an outside network.~~

~~The external network may operate either with a one-step or with a two-step clock.~~

~~The HSR node may be either slave or grandmaster. Such an HSR node should implement a Transparent Clock between the interlink port and the ring. Figure A.9 shows that the RedBox at the same time translates from 2-step to one-step clocks.~~




Figure A.9 — Attachment of a GMC to an HSR ring through a RedBox

A.6 — PRP to HSR Mapping

A.6.1 — Connection methods

A PRP network is connected to an HSR ring through two RedBoxes. There are two different ways of proceeding for connecting the networks: through Transparent Clocks and through Boundary Clocks.

A.6.2 — Transparent Clocks connection

When the two RedBoxes implement Transparent Clocks, they inject four Syncs into the ring. All four Syncs have the same sequence number, inherited from the GMC. Therefore, this method works only if the PRP LANs network elements are PRP-aware and do not remove the RCT when forwarding Sync or Announce. As in A.4.1, this rules out Boundary Clocks in the PRP LANs.

This way, the nodes can filter out the duplicates, under the assumption that the clocks have nearly the same time and that taking Syncs from different Boundary Clocks can reduce jitter.




Figure A.10 – PRP to HSR coupling by Transparent Clocks

A.6.3 PRP- HSR connection by Boundary Block

In this solution, the RedBoxes implement a Boundary Clock. In this case, the two BCs are treated as redundant clocks on HSR (see Figure A.11). The original PRP RCT is lost. Only one clock is actually sending Syncs, the other sends Announces only in case it ceases to sense the other clock. Since both BCs reference the same GMC, they are only distinguished by their sourcePortIdentity. For an HSR node, this situation is not distinguishable from a GMC in the HSR ring.




Figure A.11 — PRP to HSR coupling by BCs

IEC 1163/12

Annex A (normative)

Clocks synchronization over redundant paths in IEC 62439-3

A.1 Overview

Annex A specifies how clocks operating according to the Precision Time Protocol (PTP, IEC 61588:2009) are attached over redundant, simultaneously active paths as in PRP and HSR with no single point of failure at the network level. The use of redundant master clocks is not considered since it is already part of IEC 61588:2009 (e.g. grandmaster clusters).


This method is applied to two PTP profiles described in Annex C:

- L3E2E (Layer 3 with end-to-end delay measurement) and
- L2P2P (Layer 2 with peer-to-peer delay measurement).

Annex D contains tutorial material restricted to the aspects of IEC 61588:2009 relevant to the redundant operation in this context.

A.2 Attachment to redundant LANs by a boundary clock

An ordinary clock (OC) may be attached to redundant masters MCA and MCB by a two-port BC (see Figure A.1) which is at the same time an OC.



IEC

Figure A.1 – Doubly Attached Clock as BC (MCA is best master)

This arrangement increases the availability of both LANs in case of failure of a master. For instance in Figure A.1, if MCB fails, port B will change its state from PASSIVE_MASTER to MASTER and MCA will synchronize OC4B over the BC. If there are several BC between the two masters, only one of them will become master, according to the BMCA.

This arrangement is covered by the IEC 61588:2009 specifications. However, it requires a modification of the BMCA (see IEC 61588:2009, Figure 28) to account for the scenario in which MCA and MCB are one and the same clock. The decision for the BMCA should consider

in addition the clock quality (see Figure A.26 and A.7.3), so when port A and port B receive the same master with the same steps-removed, one of the two ports will go to PASSIVE_MASTER. This mode will not be further detailed since it is covered by IEC 61588:2009.

NOTE It has been claimed that the coupling of LAN A and LAN B by numerous BCs could introduce common failure modes.

A.3 Attachment to redundant LANs by doubly attached ordinary clocks

IEC 61588:2009 foresees that an ordinary clock has only one port. Clause A.7 extends the IEC 61588:2009 model to support an ordinary clock with redundant, paired ports, called a doubly attached clock (DAC).


The DAC principle can be applied to PRP (see Clause A.4) and to HSR (see Clause A.5). A DAC may be combined with a TC or an HC, whose closest counterpart in IEC 61588:2009 is called “*Combined ordinary and peer-to-peer transparent clock*”.

Figure A.2 shows the conceptual PTP clock model of an ordinary clock OC2 (doubly-attached slave) with two ports, one port attached to LAN A with Master Clock A (MCA) and one port attached to LAN B with Master Clock B (MCB), MCA being of better quality than both MCB and OC2. The LANs are symbolized by bridges acting as TCs.

NOTE 1 LANs can also contain BCs, in which case the BCs nearest to OC2 have the role of MCA or MCB

The redundant ports execute the PTP protocol independently. In particular, each port has a state machine according to IEC 61588:2009, 9.2.5. Each port executes the BMCA, calculates the link delays and responds to PTP management messages independently.

If a DAC has a lower clock quality than MCA or MCB, only the port connected to the better master will go to the SLAVE state and synchronize OC2, while the other port goes to the PASSIVE_SLAVE state, which does not synchronize OC2, but otherwise operates as slave port and for instance sends the delay calculation messages. The latter port will go to the SLAVE state if the other port loses synchronization or when the master on its LAN gets a better quality than the master on the other LAN.




IEC

Figure A.2 – Doubly Attached Clock when MCA is best master

NOTE 2 This concept can be extended to more than two ports or to several masters simultaneously active over the same port.

If a DAC has a better quality than both MCA and MCB, both its ports go to the MASTER state. They will send the Sync messages at approximately the same time (depending on network traffic).

Figure A.3 shows a topology in which one DAC is master (OC1) and the other is slave (OC2). Since both ports of OC2 are controlled by the same master OC1, OC2 will choose one port as SLAVE, depending on its BMCA. Since the clock identifier and quality are equal in this case, the BMCA is extended with application-specific criteria (e.g. dependent on the received clock quality) before the port identifiers are considered as a tie-breaker. A hysteresis behaviour is recommended so noisy behaviour does not cause frequent switchover.



IEC

Figure A.3 – Doubly attached clocks when OC1 is best master

NOTE 3 Such a DAC is simpler than a full BC, but this is not the main motivation. This arrangement removes the coupling between the LANs that would exist with a BC.

There are underlying dependencies between the ports. This leads to the concept of “paired ports”, which is reflected in the MIB. The dependencies are:

- The BMCA and the clock quality of OC2 is common for both ports.
- The state of one port is communicated to the other to ensure that one port does not go to the MASTER state when the other is in the SLAVE state.
- The two ports may exchange additional information, i.e. the calculated link delays.

Otherwise, when the ports are independent, the clock may send the corresponding PTP messages with the same or a different PTP sequence number on both ports.

A.4 PRP mapping to PTP

A.4.1 Scenarios and device roles

The general topology considered is shown in Figure A.4, which is not an actual topography, but only serves to explain the scenarios.

The duplicated PRP network consists of two LANs, LAN A and LAN B. The LANs are symbolized by bridges acting as TCs. BCs may exist in the LANs (not shown in Figure A.4). Ordinary clocks may be attached to the redundant LAN A and LAN B through RedBoxes as master (LAN C) and slave (LAN D).

The generalized network consists of the following elements, which are roles that a device may take, not actual devices:

- 1) OC1: DAC attached to both LANs, as master;
- 2) OC2: DAC attached to both LANs, as slave;
- 3) OC3: SAC attached to one of the PRP LANs, as master;
- 4) OC4: SAC attached to one of the PRP LANs, as slave;
- 5) OC5: SAC attached through a RedBox outside of the PRP LAN, as master;
- 6) OC6: SAC attached through a RedBox outside of the PRP LANs, as slave;
- 7) BC7: RedBox M attached to the singly attached master clock OC5, as DABC;
- 8) BC8: RedBox S attached to the singly attached slave clock OC6, as DABC;
- 9) TC9: RedBox M attached to the singly attached master clock OC5, as DATC or SLTC;
- 10) TC10: RedBox S attached to the singly attached slave clock OC6 as DATC or SLTC.

NOTE In reality, the LAN will carry a combination of all these cases, so this distinction is only conceptual.




Figure A.4 – Elements of PRP networks

A.4.2 Operation in PRP

The general PRP operating mode is that a DANP (OC1 and OC2 in Figure A.4) is attached to two distinct LANs (LAN A, LAN B in Figure A.4). For all traffic except PTP, a DANP source sends over both LANs to the same frame after appending the RCT. A frame is not modified while transiting through the LANs. The DANP destination receives the first frame of a pair and discards the duplicate on the base of its RCT and source address.


This general PRP operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same in the two LANs and some frames are modified in the TCs (bridges) while transiting through the LAN. A DANP does not receive the same PTP message from both ports even when the source is the same.

In particular:

- 1-step Syncs messages are modified by TCs to adjust the correction field.
- Follow_Up frames are generated by every 2-step clock and carry no RCT.
- TCs are not PRP-aware and not obliged to forward the RCT, which is a message part that comes after the payload.
- BCs present in the LAN are not PRP-aware and would generate their own Announce and Sync frames with no RCT appended.

Figure A.5 shows the connection between the master OC1 and a slave OC2 in a PRP network, both devices being DANPs (and DACs). The path of a Sync message through the LANs with the correction of the clock for the residence delay and link delays is suggested using a 2-step clock as an example. OC2 receives from each port the same PTP messages coming from the same master with different delays.

Since the RCT of PRP cannot be relied upon for PTP messages, ordinary clocks will not perform duplicate discard on them and will consider the two ports as independent.



IEC

Figure A.5 – Connection of a master clock to an ordinary clock over PRP

A.4.3 Configuration specification

When operating on a PRP network, an ordinary clock shall be either:

- 1) a DAC, attached directly as a DANP (see A.4.4, e.g. OC1 or OC2 in Figure A.4), in which case the clock obeys to the IEC 61588:2009 default BMCA extended by this specification;
- 2) a SAC, which is unaware of this specification, but is supposed to obey to the default BMCA IEC 61588:2009, attached either:

- to one LAN only as a singly attached clock (see OC3, OC4 in Figure A.4) without the benefit of redundancy;
- through a RedBox as specified in A.4.5 (see OC5, OC6 in Figure A.4).

NOTE A PRP topology in which one MC is located on LAN A and another MC on LAN B is considered a degraded situation, since a grandmaster clock (GMC) on one LAN ignores another master on the other LAN and cannot serve as a back-up. This situation happens when links are lost and ordinary clocks take over as master for lack of a better alternative. This situation can be avoided if at least one device according to A.7.3. exists. However, this can also be an intentional way of providing independent redundant clocks.

A.4.4 Specifications of DANP as DAC

A.4.4.1 Sending in the DACs

A DANP shall send the PTP messages over both ports according to the port state of each port, as selected by A.7.4.

NOTE There is no need to append an RCT but if one is added it would be ignored.

A.4.4.2 Reception in the DACs

A DAC shall handle all PTP messages according to A.3 for each port separately;

NOTE 1 This means in particular that a port exploits the Sync and Delay_Resp (in E2E) for supervision even if it does not use them for synchronization.

NOTE 2 A DANP does not rely on the RCT of received PTP messages.

A.4.5 Clock model of a RedBox for PTP

A.4.5.1 PRP RedBox types

RedBoxes connect singly attached devices to a duplicated LAN. There exists different kinds of RedBoxes with regard to the redundant operation. All kinds can be used and mixed in the same network. Some RedBox types introduce constraints that affect the entire network.


The following types of RedBox are considered:

- 1) RedBoxes as three-port BCs (TWBC)
The RedBox operates as a three-port BC, which acts as a slave on exactly one LAN and as a master or passive on the other two LANs. This increases resiliency of singly attached clocks on LAN A or LAN B since the BC can bridge the two LANs. This scenario is fully covered by the present IEC 61588:2009 specifications and is not repeated here (see A.4.5.2).
- 2) RedBox as a doubly attached BC (DABC)
The RedBox operates as a Doubly Attached Clock towards the duplicated LAN and as a BC towards the single LAN. The time distribution principles may be different on both sides of the BC; there is no restriction of IEC 61588:2009. (see A.4.5.3)
- 3) RedBox as a doubly-attached TC (DATC)
The RedBox operates as a DAC towards the duplicated LAN and as a TC towards the single LAN and for this it uses the BMCA to choose one path. They route the PTP messages to the correct path by manipulating the source port identity on the single-thread LAN. Restrictions to the IEC 61588:2009 are that a singly attached master receives twice as many Delay_Req as in the non-redundant case and that the number of ports in a device is restricted to 4 096 (see A.4.5.4)
- 4) RedBox as a stateless TC (SLTC)
This RedBox operates in cut-through mode and can only discard messages at reception. These RedBoxes provide higher performance for all messages. The SLTC tags PTP messages on all LANs.
The Singly Attached Clocks on the single-threaded LAN are able to operate with two simultaneous masters, as specified in IEC 61588:2009.
An SAC master on a single-threaded LAN receives twice as many Delay_Req as it would

without redundancy, causing additional processor load, and a total of four Delay_Respon would circulate on the LANs, causing additional traffic. The number of ports in a device is restricted to 4 096 (see A.4.5.5)

A.4.5.2 PRP RedBox as three-way BC (TWBC)

RedBoxes can be implemented as three-way BCs, (see A.2 and Figure A.6). This case is covered by the IEC 61588:2009 specifications.



IEC


Figure A.6 – PRP RedBox as BCs (OC3 and BC7 are best masters)

A.4.5.3 PRP RedBox as doubly attached BC (DABC)

A.4.5.3.1 RedBox as DABC – principle

A RedBox operates as a BC with two ports, one virtual port consisting of two physical ports, A and B operated as DAC towards the duplicated LAN A and LAN B and one port towards the single path LAN C. It does not operate as a BC between LAN A and LAN B.

Figure A.7 shows the conceptual PTP clock model of a RedBox as a DABC.



IEC

Figure A.7 – RedBox DABC clock model

In this configuration, two cases are distinguished:

- 1) The best master resides on the PRP LANs (OC1 in Figure A.7), port A or port B become slaves and port C competes for mastership on the non-PRP side. If it wins the BMC competition, the master OC1 will synchronize the BC, which will in turn synchronize OC6.
- 2) The best master resides on the single-threaded LAN (e.g. OC5 in Figure A.8), port C becomes slave and synchronizes the BC, whose port A and port B compete for mastership on their respective LANs, and, if successful, synchronize the DAC OC2.

A.4.5.3.2 PRP RedBox as DABC with E2E

A.4.5.3.2.1 PRP RedBox as DABC with E2E – principle

When a RedBox implements a BC with E2E (see Figure A.8), the exchange of PTP messages is limited to the individual LANs. Ordinary clocks attached to LAN A and LAN B obey to the doubly-attached clock model; they implement the E2E specified in IEC 61588:2009.

Clocks in LAN C and LAN D can use another link delay calculation (e.g. clock OC6 in Figure A.8, which uses P2P) or even a different time distribution protocol such as SNTP.




Figure A.8 – PRP RedBoxes as DABC with E2E – BC7 is master

A.4.5.3.2.2 PRP RedBox as DABC with E2E – timing diagram

Figure A.9 shows the timing of the PTP messages between the two RedBoxes. The RedBoxes behave like DACs. The PTP message exchange on the single-threaded LANs is shown in Figure A.9 with E2E as example.





Figure A.9 – PRP RedBoxes as DABC with E2E – timing

A.4.5.3.3 PRP RedBox as DABC with P2P

A.4.5.3.3.1 PRP RedBox as DABC with P2P – principle

When a RedBox implements a BC with P2P (see Figure A.10) the exchange of Sync messages is limited to the individual LANs. The P2P is link-local. The clocks attached to both LAN A and LAN B obey to the Doubly Attached Clock rules.

The time distribution protocol on LAN C or LAN D can be different, Figure A.10 shows an example of P2P on LAN C and E2E on LAN D.




IEC

Figure A.10 – PRP RedBoxes as DABC with P2P – OC5 is best master

A.4.5.3.3.2 PRP RedBox as DABC with P2P – timing diagram

Figure A.11 shows the timing diagram of the P2P, 1-step PTP message exchange between the RedBoxes. The BCs implement E2E on the single-threaded LANs to demonstrate the independence of the protocols.



IEC

Figure A.11 – PRP RedBoxes as DABC with P2P – timing

A.4.5.3.4 PRP RedBox as DABC – specification

RedBoxes acting as BCs shall behave as Doubly Attached Clocks according to A.4.4 on the PRP side.

NOTE This specification applies both to P2P and E2E. The behaviour on the single-threaded side is not specified, it is not constrained to be PTP.

A.4.5.4 PRP RedBox as Doubly Attached TC (DATC)

A.4.5.4.1 PRP RedBox as DATC with E2E

A.4.5.4.1.1 PRP RedBox as DATC with E2E – principle

All messages needed for delay measurement need to take the same path:
in 1-step, E2E, Sync and Delay_Req need to take the same path, while
in 2-step, E2E, Sync, Follow_Up, Delay_Req and Delay_Resp need to take the same path.

A DATC ensures this by letting the RedBox execute the same BMCA algorithm as a DAC would.

When the ordinary clock is a DAC (OC1, OC2), each port keeps track of the port over which it sends and receives the PTP messages, LAN A and LAN B are separate.

The most complex situation is when both the master and the slave clock are SAN attached through a RedBox, (master OC5 or slave OC6, as shown in Figure A.12). These devices

ignore over which LAN A or LAN B the PTP messages arrived. In this case, the RedBox S (TC10) makes the choice.




Figure A.12 – PRP RedBox as DATC with E2E –signal flow

The RedBox M on the master side (TC9 in Figure A.12) duplicates the Sync coming from the master clock (OC5) over LAN_C and forwards it to both port A and port B, with individual correction of the residence time. The Sync messages transit over LAN A and LAN B with different delays.

Two cases are considered (see Figure A.12):

- 1) When the slave clock is a DAC that operates with PRP (OC2), it receives both Sync and accepts one according to its BMCA. OC2 sends two Delay_Req on LAN_A and LAN_B that travel to the master OC5. The DAC OC2 will only use the Delay_Resp for time correction on the same LAN as it accepted the Sync, but the port in the PASSIVE_SLAVE state will nevertheless be able to calculate the delay on the other part, so switchover can be smoother.

- 2) When the slave clock is a SAC (OC6) that resides behind a RedBox S (TC10 in Figure A.12), TC10 receives two Syncs and selects the one from the better master according to its BMCA (A.7.4) and forwards the Sync to the slave clock (OC6) on LAN C. The slave clock OC6 responds with one Delay_Req. RedBox S (TC10) duplicates the received Delay_Req on both LANs.

The master OC5 will respond to each Delay_Req with a Delay_Resp. To ensure that the Delay_Resp takes the same path as the Delay_Req, the DATC RedBox M (TC9) manipulates the port identity of the Delay_Req, by setting the most significant two bits Bit 15 and Bit 14 to “10” for LAN A and to “11” for LAN B respectively, before forwarding it on LAN C to OC5.

The master OC5 copies the port number of the sourcePortIdentity in the Delay_Req to the port number of the requestingPortIdentity in the Delay_Resp. When the Delay_Resp returns from the master OC5, the RedBox M routes the Delay_Resp according to the manipulated port number on the same path as the Delay_Req came from. At the same time, the RedBox M sets the two most significant bits Bit 15 and Bit 14 of the port number in the requestingPortIdentity field to “00” to ensure that the slave clock recognizes its identity.

RedBox S (TC9 in Figure A.12) accepts the Delay_Resp only from the LAN over which it received the Sync (and Follow_Up in 2-step) and forwards it over LAN D; it rejects the other Delay_Resp, regardless which one came first, but registers its timely arrival to supervise the unused path in the event of a unidirectional failure.

NOTE 1 This assumes that the portNumber of a DAC on LAN A and LAN B is “00” in the two most significant bits.

NOTE 2 The RedBox also sends Delay_Req on the inactive LAN to supervise that both LANs are operating. The cost is that a master on LAN C receives twice as many Delay_Req as in the non-redundant case.

NOTE 3 This behaviour is not strictly that of a TC according to IEC 61588:2009, which does not foresee that TCs perform the BMCA. However, it is the same as for the DAC and for the BCs.

A.4.5.4.1.2 PRP RedBox as DATC with E2E –timing

The timing diagram in Figure A.13 shows the 1-step PTP message exchange between the master (OC5 in Figure A.12) and the slave (OC6 in Figure A.12) over the PRP LAN A and LAN B.

NOTE The 2-step exchange looks similar, only that a Follow_Up follows the Sync.




Figure A.13 – PRP RedBox as DATC with E2E – timing

The RedBox M on the master side (TC9 in Figure A.12) duplicates the Sync coming from the master clock (OC5) over LAN_C and forwards it to both port A and port B, with individual correction of the residence time. The Sync messages transit over LAN A and LAN B with different delays

Assuming that the RedBox S on the slave side (TC10 in Figure A.12) considers LAN A as carrying the better master, it forwards the Sync from LAN A to LAN D, correcting its residence time. It also supervises that Sync B arrives to supervise the passive port.

When the corresponding Delay_Req arrives (same sequence number as Sync), RedBox S (TC10 in Figure A.12) forwards it only on the LAN from which it propagated the Sync to the slave.

It may send the Delay_Res over the inactive channel to supervise that no unidirectional failure exists on the path.

The RedBox M (TC9 in Figure A.12) forwards the (both) Delay_Req to the master (OC5 in Figure A.12).


The master returns the Delay_Res and RedBox M (TC9 in Figure A.12), forwards this response on both LANs (since it cannot know which port RedBox S prefers).

RedBox S (TC9 in Figure A.12) accepts the Delay_Resp only from the LAN over which it received the Sync (and Follow_Up in 2-step) and forwards it over LAN D; it rejects the other Delay_Resp, regardless which one came first, but registers its timely arrival to supervise the unused path in the event of a unidirectional failure.

A.4.5.4.2 PRP RedBox as DATC with P2P

A.4.5.4.2.1 PRP RedBox as DATC with P2P – principle

A RedBox can be implemented as a DATC as the model of Figure A.14 shows.



IEC

Figure A.14 – PRP RedBox as DATC with P2P

The RedBox M on the master side (TC9 in Figure A.14) duplicates the Sync.

The RedBox S on the slave side (TC10 in Figure A.14) only forwards one of the Sync.

In P2P, it does not matter if the Sync travels over LAN A and LAN B: the corrections are per TC. But in 2-step, the Follow_Up needs to take the same way as the Sync. Some clock profiles require that Announce takes the same path as the Sync.

The Sync that is not forwarded by RedBox S is nevertheless used as a health check for the path.

A.4.5.4.2.2 PRP RedBox as DATC with P2P – timing

Figure A.15 shows the timing diagram of the RedBoxes as DATCs using P2P. LAN A is assumed to contain the preferred master. Sync B is transmitted as a health check of PTP synchronization on LAN B.




Figure A.15 – PRP RedBox as DATC with P2P – timing

NOTE The timing diagram for 2-step is similar, except for the Follow_Up that follow the Sync and the Pdelay_Resp_Follow_Up that follow the Pdelay_Resp.

A.4.5.4.3 PRP RedBox as DATC – specification

A RedBox operating as a DATC shall act as a TC according to IEC 61588:2009, operate as DAC towards the PRP network and implement the paired port selection of A.7.4.

A RedBox operating as a DATC, when in E2E mode shall:

- when it receives from LAN A or LAN B a PTP message, except Delay_Req, forward this message to the single-threaded LAN only if it came from the LAN with the best master;
- when it receives from LAN A or LAN B a Delay_Req message, set the most significant two bits (MSB) of the sourcePortIdentity.portNumber of the Delay_Req message to “10” for LAN A and to “11” for LAN B respectively;
- when it receives from the single-threaded LAN a PTP message except Delay_Resp, forward a copy to both LAN A and LAN B;
- when it receives from the single threaded LAN a Delay_Resp, inspect its requestingPortIdentity.portNumber, restore it to the value it had in sourcePortIdentity.portNumber of Delay_Req and forward it only to the LAN where the Delay_Req came from.

A RedBox operating as a DATC when in P2P mode shall:

- forward a PTP message that it receives from LAN A or LAN B to the single-threaded LAN, but only if it came from the LAN with the best master and if it is neither a Delay_Req nor a Delay_Resp message;
- forward a copy of a PTP message that it receives from the single-threaded LAN to both LAN A and LAN B, except for Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req;
- send and receive the Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req messages and calculate the link delays as specified in IEC 61588:2009 on all ports individually.

A.4.5.5 PRP RedBox as stateless TC (SLTC)

A.4.5.5.1 PRP RedBox as SLTC – principle

In applications where the RedBox cannot implement the BMCA (e.g. simple hardware cores which forward frames in cut-through mode), the RedBox can be constructed as a stateless TC which does not discard PTP duplicates, but just sends all PTP messages it receives from LAN A or LAN B on the single-thread LAN C or LAN D, tagging them as coming from LAN A or LAN B by manipulating the bit 13 and bit 12 of the source port identity.

Singly attached clocks on the single-thread LAN C and LAN D will thus receive all PTP messages twice. Therefore, all OCs (e.g. OC6) execute the BMCA to pick up the correct Sync and Announce.

A.4.5.5.2 PRP RedBox as SLTC – limitations

This type of RedBox operates with E2E or P2P, 1-step or 2-step and over multiple domains. It expects that the OCs on the single-thread LAN C or LAN D follow the IEC 61588:2009 standard, and in particular the state machine in IEC 61588:2009, Figure 24, which prescribes that the slave clock only listens to the best master in case it receives Syncs from several masters.

This type of RedBox reserves the four most significant bits of the source port identity to tag the frames, therefore the number of ports per device is reduced to 4 096. This affects also clocks that are not redundancy-aware.

This method is vulnerable to a unidirectional link failure in the RedBox S to RedBox M direction because, if Delay_Req is lost Delay_Resp will not be received.

With this method, slave devices will report a portIdentity of the selected master which does not actually exist; this might cause some confusion to network administration.

An SLTC cannot be used to connect two duplicated HSR or PRP networks.

NOTE This method is therefore not applicable to slave-only clocks which do not execute the BMCA.

A.4.5.5.3 PRP RedBox as SLTC with E2E

Figure A.16 show the most complicated case in which the master OC5 sends an Announce or a Sync, which is duplicated by RedBox M and travels on LAN A and LAN B. When each of the two Announce or Syncs arrives at RedBox S (TC10), it tags them as coming from LAN A or LAN B by manipulating the source portNumber bit 13 and bit 12 and forwards them on LAN D.

OC6 receives two Announces and selects one of them as its best master. The different port identities serve as a tie-breaker. OC6 will from now on receive Syncs only from the selected master.

To measure link delay in end-to-end, OC6 sends a Delay_Req with its source port identity, which RedBox S receives. RedBox S manipulates the bit 13 and bit 12 of the source port identity when sending them over LAN A (value 10) and for LAN B (value 11).

NOTE 1 The reason for tagging the Delay_Req at RedBox S rather than at RedBox M as for a DATC is that an SLTC RedBox M cannot modify the tags since the SLTC RedBox S needs them.

Both Delay_Req travel to RedBox M, which sends them both unchanged to the master OC5. The master OC5 responds to each of them with a Delay_Resp.

NOTE 2 Master OC5 can notice that it receives redundant Delay_Resp since both come from a device with the same clockIdentity.


The RedBox M duplicates each Delay_Resp on both LANs, without modifying them. In a fault-free network four Delay_Resp are under way.

RedBox S can remove two of the Delay_Resp from the network, those where the tagging indicates that Delay_Req was on the other LAN (requestingPortIdentity not corresponding to the LAN the frame was received from), cases “10” and “01” in Figure A.17.

NOTE 3 If the RedBox cannot suppress the sending of the message, it can alternatively leave the requesting port identity untouched so the slave will not recognize this as a response to his request. Alternatively, it could also corrupt the CRC so the next network node will discard the frame.

RedBox S tags the Delay_Resp as it would tag the Syncs. Therefore, the OC will only take one and discard the other.

RedBox S resets the sets bit13 and bit 12 in the requesting port identity in the Delay_Resp back to “00” when sending the response to LAN D so the slave OC6 identifies this as his Delay_Resp to his Delay_Req. RedBox S does this only if the requesting port identity tagging in the Delay_Resp matches the source port identity tagging in the Delay_Req. Otherwise it leaves the requesting port identity untouched.



IEC

Figure A.16 – PRP RedBox as SLTC with E2E

The timing diagram in Figure A.17 shows the message exchange.





Figure A.17 – PRP RedBox as SLTC with E2E – timing

RedBox M can also operate as a DATC. In this case, it will send a Delay_Resp on the same LAN as it received the Delay_Req. It will then only send two Delay_Resp, one on each LAN A and LAN B, which will conserve the sourcePortIdentity.portNumber of Delay_Req in their requestingPortIdentity.PortNumber.

NOTE This was not necessary in a DATC, since a DATC does not operate in cut-through mode and can therefore select the messages it sends.

A.4.5.5.4 PRP RedBox as SLTC with P2P

The operation is similar to A.4.5.5.3, except that there are no Delay_Req and Delay_Resp messages (see Figure A.18). The peer delay messages are link-local and therefore unaffected. OC6 receives two Sync with a difference in the upper nibble and selects the best one as master.



IEC

Figure A.18 – PRP RedBox as SLTC with P2P

A.4.5.5.5 PRP RedBox as SLTC – specifications

A RedBox operating as SLTC shall operate as TC according to IEC 61588:2009.

A RedBox operating as SLTC, when operating in E2E, shall:

- if it receives a PTP message from a single-threaded LAN,
 - duplicate this message to LAN A and LAN B, and if the message is a Delay_Req, set the bit 13 and bit 12 in the most significant nibble (MSN) of the sourcePortIdentity.portNumber to “10” if it forwards it to LAN A and to “11” if it forwards it to LAN B, respectively.
- if it receives a PTP message from LAN A or LAN B,

- if it is a Delay_Resp message, discard it, if bit 13 and bit 12 of the requestingPortIdentity.portNumber do not match the value which the RedBox set in sourcePortIdentity.portNumber of Delay_Req on the LAN and if they match, reset these two bits to “00” before forwarding the Delay_Resp to the single-threaded LAN;
- otherwise, if it receives a PTP message from LAN A or LAN B in which the sourcePortIdentity bit 13 and bit 12 are 00, set them to “10” if it received it from LAN A, and to “11” if it received it from LAN B respectively, and forward the PTP message to the single-threaded LAN with possible correction.

A RedBox operating as SLTC, when operating in P2P, shall:

- forward a PTP message that it receives from LAN A or LAN B to the single-threaded LAN and if it is neither a Delay_Req nor a Delay_Resp message;
- forward a copy of a PTP message that it receives from the single-threaded LAN to both LAN A and LAN B, except for Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req;
- Send and receive the Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req messages and calculate the link delays as specified in IEC 61588:2009 on all ports individually.

A.5 HSR Mapping to PTP

A.5.1 PTP traffic in HSR

The operating mode of HSR is that a DANH sends the same frame in both directions into the ring, prefixed with the HSR Tag containing a sequence number. Each DANH relays the frame until it reaches the source again or is otherwise discarded. A destination DANH therefore receives two frames, uses the first and discards the duplicate.

This operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same going clockwise or counter-clockwise in the ring. A DANH does not receive the same message from both ports since some messages are modified by intermediate nodes. Figure A.19 illustrates the paths of the different PTP messages in an HSR ring network.





Figure A.19 – HSR with one GMC


For 1-step operation, each DANH implements the 1-step Sync and P2P and exchanges the messages shown in Figure A.20.



IEC

Figure A.20 – PTP messages sent and received by an HSR node (1-step).

Alternatively, the HSR nodes may operate with 2-step correction as Figure A.21 shows.



IEC

Figure A.21 – PTP messages sent and received by an HSR node (2-step)

A.5.2 HSR nodes specifications

A DANH shall implement a Hybrid Clock, made of a TC and an optional OC.

A DANH shall send all PTP messages with a HSR tag.

NOTE 1 HSR tagging is in principle not needed for Pdelay_Req and Pdelay_Resp, which are link-local, but it allows to remove corrupted frames out of the ring, although in principle an HSR node does not forward untagged frames.

A DANH acting as master clock shall send all PTP messages in both directions with an HSR tag and with the same PTP sequence number.

A DANH shall handle all PTP messages for each port separately.

A DANH shall send and receive the Pdelay_Req and Pdelay_Resp messages per port and compute the link delay per port.

A DANH receives from each port one Announce and one Sync message coming from the same master with the same Clock Identity and different delays. It shall forward both in the same direction with the appropriate correction, except if this contradicts the HSR rules.

A clock shall apply the Best Master Clock algorithm on port A and port B, taking its ordinary clock as the common reference.

NOTE 1 TC-only DANH have been considered, but this would have added an unnecessary option.

If the Announce and Sync messages from both ports have the same Clock Identity, a DANP may use Syncs from both ports for synchronization/syntonization, or from one port only, a DANH shall listen to the port that provides the best clock quality.

A port that receives 2-step Sync message shall await the Follow_Up message from the same port before using the Sync timestamps to adjust its clock or to perform 2-step to 1-step conversion.

All PTP messages except the link-local Pdelay_Req/Pdelay_Resp shall use as source address the MAC address of the node that injected the frame into the ring. The link-local Pdelay_Req and Pdelay_Resp shall insert the MAC address of the TC.

NOTE 2 The IEEE P1588 revision group foresees that all devices that modify a frame insert their own MAC address as source address. Doing so would prevent removing circulating messages from the ring.

A.5.3 Redundant clocks in HSR

As Figure A.19 shows, there can be a redundant master clock in the ring. However, both cannot be active at the same time in PTP. Therefore, the grandmaster clock is selected among the master-capable clocks according to the Best Master Clock of PTP.

In case of loss of the master, the back-up master detects the loss and start broadcasting its own Announce messages.


A.5.4 Attachment of an MC to an external LAN

An MC may be connected to an HSR ring through a RedBox (see Figure A.22).

In its simplest form, this RedBox can be a three-port HSR node connected through an interlink port to an outside network.

The external network may operate either with a 1-step or with a 2-step clock.

The HSR node may be either slave or master. Such an HSR node may implement a TC or a BC between the interlink port and the ring. Figure A.22 shows that the RedBox at the same time translates from 2-step to 1-step clocks.



IEC

Figure A.22 – Attachment of a GMC to an HSR ring through a RedBox as TC

This arrangement allows to use an existing DANH to inject the synchronization. However it is not recommended when the MC is to be attached redundantly.

A.6 PRP to HSR Mapping

A.6.1 Connection methods

A PRP network is connected to an HSR ring through two RedBoxes. There are two different ways of proceeding for connecting the networks: through BCs (recommended) and through TCs (not recommended).

A.6.2 PRP-HSR connection by BC

The RedBoxes implement each a BC. The two BCs are treated as redundant clocks on HSR (see Figure A.23). The original PRP RCT is lost. Only one RedBox is actually sending Announces and Syncs, the other RedBox starts to send Announces and Syncs only in case it ceases to sense the other clock, but its port in the HSR ring remains otherwise in the PASSIVE_MASTER state (i.e. the second RedBox is not synchronized by the HSR ring).

Since both BCs reference the same GMC, they are only distinguished by their sourcePortIdentity.




Figure A.23 – PRP to HSR coupling by BCs

A.6.3 PRP-HSR connection by TCs

The redundant PRP-HSR connection by a TC causes the injection of four PTP messages into the ring, see Figure A.24. This configuration is not recommended and not further specified.





Figure A.24 – PRP to HSR coupling by TCs

A.7 Doubly attached clock model

A.7.1 State machine

Each port can be in the port states defined in IEC 61588:2009, Table 8, as Figure A.25 shows. The transitions that have been added to the state machines given in IEC 61588:2009, 9.2 for the Doubly Attached Clock are in red.



IEC

Figure A.25 – Port states including transitions for redundant operation

Table A.1 shows the states as defined in IEC 61588:2009, 9.2.5, Table 10 with additions (underscore) and deletions (strike-through).

Table A.1 – States

State	Definition
INITIALIZING	While a port is in the INITIALIZING state, the port initializes its data sets, hardware, and communication facilities. No port of the clock places any PTP messages on its communication path. If one port of a BC is in the INITIALIZING state, then all ports go to the INITIALIZING state.
FAULTY	The fault state of the protocol. A port in this state places no PTP messages except for management messages that are a required response to another management message on its communication path. In a BC, no activity on a faulty port affects the other ports of the device. If fault activity on a port in this state cannot be confined to the faulty port, then all ports go to the FAULTY state.
DISABLED	The port places no messages on its communication path. In a BC, no activity at the port affects the activity at any other port of the BC. A port in this state discards all PTP received messages except for management messages.
LISTENING	The port is waiting for the announceReceiptTimeout to expire or to receive an Announce message from a master. The purpose of this state is to allow orderly addition of clocks to a domain. A port in this state places no PTP messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, or signaling messages, or management messages that are a required response to another management message.
PRE_MASTER	The port behaves in all respects as though it were in the MASTER state except that it places no messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, signaling, or management messages.
MASTER	The port behaves as a master port.
PASSIVE_MASTER	The port places no messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, or signaling messages, or management messages that are a required response to another management message.
UNCALIBRATED	One or more master ports have been detected in the domain. The appropriate master port has been selected, and the local port is preparing to synchronize to the selected master port. This is a transient state to allow initialization of synchronization servos, updating of data sets when a new master port has been selected, and other implementation-specific activity.
SLAVE	The port is synchronizing to the selected master port.
PASSIVE SLAVE	The port is not synchronizing to the selected master port. In E2E, it may send Delay_Req; or in P2P, it sends Pdelay_Req, Pdelay_Resp, (Pdelay_Resp_Follow_Up). It may send signaling or management messages that are a required response to another management message.
NOTE “PASSIVE_MASTER” and “PASSIVE_SLAVE” are jointly expressed as “PASSIVE” in the management interface	

Table A.2 lists the transitions that cause state changes. In particular, the transition BMC_PSLAVE has been introduced to keep the paired ports synchronized.

Table A.2 – Transitions

Transition	Definition
ANNOUNCE_TO	No Announce has been detected on this port within the timeout. (the other port may still receive Announce messages)
BMC_SLAVE	Outcome of STATE_DECISION_EVENT: this port receives the best master. (the other port will go to PASSIVE_SLAVE)
BMC_MASTER	Outcome of STATE_DECISION_EVENT: the local OC is the best clock on both ports (the other port will also go to MASTER).
BMC_PASSIVE	Outcome of STATE_DECISION_EVENT: this OC is not the best but the port is prevented from entering the MASTER state (e.g. it is a grandmaster that cannot accept synchronization).
BMC_PSLAVE	Outcome of STATE_DECISION_EVENT: the clock on this port has a worse quality than that on the other port (the other port is in the SLAVE state).
DESIGNATED_DISABLED	A command has been given to disable this port.
DESIGNATED_ENABLED	A command has been given to enable this port.
FAULT_DETECTED	An error occurred. The logic is supposed to do some self-checking.
FAULT_CLEARED	The error has been corrected.
INITIALIZE	Reception of an INITIALIZE management message.
MASTER_CLOCK_SELECTED	The port has synchronized to the master on this port.
POWER_UP	The port starts operation.
QUALIFICATION_TO	The qualification time out (see IEC 61588:2009, 9.2.6.10) expired.
Ready	The initialization procedure is finished.
STATE_DECISION_EVENT	Outcome: BMC_MASTER, BMC_SLAVE, BMC_PSLAVE, BMC_PASSIVE
SYNC_FAULT	The clock synchronization has performed a correction outside of the allowed range.

Table A.3 indicates the auxiliary variables used in the state chart of Figure A.25.

Table A.3 – Variables

Variable	Definition
new_master	identity of the higher quality master
old_master	identity of the current master

A.7.2 Supervision of the port

A port in the SLAVE or PASSIVE_SLAVE state enters the FAULTY state when it ceases to receive the expected PTP messages. This governs the FAULTY_SYNC and the STATE_DECISION_EVENT transition.

To this effect, the ports shall implement a time-out on the PTP messages: Sync (and Follow_Up), Delay_Resp (in end-to-end) and Pdelay_Resp (in peer-to-peer).

The time-out on these messages shall be set by default to the double value of the expected period (logMessageInterval) plus 60 %.

NOTE 1 IEC 61588:2009 has no time-out on the PTP messages except Announce. This could change in a new edition.

NOTE 2 Only one outage is tolerated before the clock switches to the redundant path. If the Sync interval is 1 s, the time-out will be raised 2,6 s after reception of a correct Sync. This accounts for the ± 30 % jitter permitted by IEC 61588:2009, 7.7.2.1.

A.7.3 BMCA for paired ports

The BMCA algorithm in IEC 61588:2009 is extended for the purpose of paired ports synchronized by the same master. If BCs exist, the Steps Removed (number of hops) can be different and serve as a selection criteria.

The reception of identical messages is a normal case and therefore an application-specific quality chooses the clock to listen to, e.g. smaller correction field (after proper filtering) or an anomaly in the correction field.

The application-specific quality is reflected in the decision “*compare quality*” in Figure A.26, which has been adapted from Figure 28 of IEC 61588:2009.




Figure A.26 – BMCA for redundant masters

This paired port BMCA specification is considered in the term “better” transition in the sequel.

A.7.4 Selection of the port state

The paired port BMCA shall define which clock synchronizes which other clock.

- a) If the OC is the best (or only) master both on port A and port B, both ports go to the MASTER state (BMC_MASTER).
- b) If the master MCA on port A or the master MCB on port B is better than the OC, the port A with the better master goes to the SLAVE state (BMC_SLAVE) while the other goes to the PASSIVE_SLAVE state (BMC_PSLAVE).
- c) If a port in the SLAVE state loses its master (ANNOUNCE_TO, SYNC_FAULT), or the BMCA detects a better master on the paired port, the port transits from the SLAVE to the PASSIVE_SLAVE state (BMC_PSLAVE), while the other port wakes up from the PASSIVE_SLAVE state and becomes SLAVE (BMC_SLAVE).
- d) If the master on one port is worse than the OC or the master on the other port, no effort is spent in trying to supply the synchronisation, this is not a BC.

A.8 PTP datasets for high availability

A.8.1 General

The objects in the IEC 61588 datasets are extended. These objects are exposed in the MIB for local or remote management. In case of discrepancy, the wording in the MIB takes precedence.

A.8.2 Data types

A.8.2.1.1 TimeInterval

This type is defined in IEC 61588, 5.3.2

A.8.2.1.2 EUI64

Base type: OCTET STRING (SIZE(8))

Description: This type is defined in IEC 61588, 7.5.2.2.2

A.8.2.1.3 PortStateType

This type is defined in IEC 61588:2009, 8.2.5.3.1, Table 8 and extended by value 10: PASSIVE_SLAVE.

A.8.2.1.4 ProfileSet

Base type: Unsigned32 (bit set)

References: IEC 62439-3:2015

Description: Bits are numbered from 0 (least significant) to 31 (most significant). Bit “7” is set for the L3E2E profile (see Clause C.5), Bit “8” for the L2P2P profile (see Clause C.6) and Bit “9” for the IEC/IEEE 61850-9-3 profile (see Annex B). Other bits are not defined. If the value of this object is “0”, no profile is defined.

A.8.2.1.5 PortAttachmentType

Base type: Enum32

Values:

- 0: not specified;
- 1: the port belongs to a singly attached ordinary clock;
- 2: the port belongs to a singly attached boundary clock;
- 3: the port belongs to a singly attached transparent clock;
- 4: the port belongs to a boundary node;
- 5: the port belongs to a doubly-attached ordinary clock (portDs.pairedPort is not >FFFF);
- 6: the port belongs to a doubly-attached boundary clock (portDs.pairedPort is not >FFFF);
- 7: the port belongs to a doubly-attached transparent clock (portDs.pairedPort is not >FFFF);
- 8: the port belongs to a stateless transparent clock (portDs.pairedPort is not >FFFF);
- 255: unknown.

A.8.3 Datasets for ordinary or boundary clocks

A.8.3.1 Additional or modified members

The members additional or modified with respect to IEC 61588 are as follows:

- defaultDs.numberPorts
- defaultDs.profileSet;
- defaultDs.timeInaccuracy;
- portDs.portState;
- portDs.portEnabled;
- portDs.dlyAssymetry;
- defaultDs.profileId;
- portDs.vlanEnable;
- portDs.vlanId;
- portDs.vlanPrio;
- portDs.twoStepFlag;
- portDs.peerIdentity;
- portDs.prpPairedPort;
- portDs.prpAttachment;
- portDs.errorCounter;
- portDs.peerDelayMax.

A.8.3.2 Dataset DefaultDS

A.8.3.2.1 Object defaultDs.numberPorts

Type: Unsigned 8 [5 to 254]

Reference: IEC 61588:2009, 8.2.1.2.3, IEC 62439-3:2015, Clause A.3

Description: The number of PTP ports on the device. (1 if OC, >1 for Boundary Clock). IEC 62439-3 allows 2 for OC that are doubly attached.

A.8.3.2.2 Object defaultDs.profileSet

Type: ProfileSet

Description: The set of profiles which the clock has the ability to support with the required performance. Profile activation is a port property.

Default value: 0

NOTE This object serves management to check that all devices are properly configured for deployment or maintenance. This information is not used in the PTP protocol. It does not guarantee that all parameters of the clock have been properly set. A clock can support several profiles, but a port supports only one.

A.8.3.2.3 Object defaultDs.timeInaccuracy

Type: TimeInterval

Units: scaledNanoseconds

Reference: IEC 61588:2009, 5.3.2, IEC 62439-3:2015, C.7.4.1 and C.7.6.1.

Description: The absolute value of the rated worst-case time error with respect to absolute time when the clock is connected to its reference signal. If the clock is a boundary clock, the reference signal is the clock signal applied to any port, and the rated time error the error observed on any master port.

Default value: 0xFFFFFFFF

NOTE This object serves to check that a MC or BC has a sufficient quality for the given application. This value is not measured at run-time, it is a design parameter that also appears in the MIB. This value is not transmitted using in the PTP protocol. The inaccuracy at run-time can be deduced from the defaultDS.clockQuality.offsetScaledLogVariance object.

A.8.3.2.4 Object defaultDs.offsetFromMasterLim

Type: Unsigned32

Units: nanoseconds

Reference: IEC 62439-3:2015, A.8.3.2.4

Description: The value of the limit to the offset from master that signals an abnormal condition.

Default value: 0xFFFFFFFF

A.8.3.3 Dataset ParentDS

As in IEC 61588.

A.8.3.4 Dataset PortDS

A.8.3.4.1 Object portDs.portState

Type: PortStateType

Reference: IEC 61588:2009, 8.2.5.3.1, Table 8; IEC 62439-3:2015, Table A.1

Description: The current state of the protocol engine associated with this port taken from the enumeration in Table 8, extended by the PASSIVE_SLAVE state in doubly attached clocks.

A.8.3.4.2 Object portDs.portEnabled

Type: Boolean

Reference: IEC 61588:2009, 8.2.5.3.1

Description: Enables the port for PTP traffic and indicates the enabled status of the port.

Default value: true

A.8.3.4.3 Object portDs.dlyAsymmetry

Type: TimeInterval

Units: scaledNanoseconds – IEC 61588, 5.3.2

Reference: IEC 61588:2009, 7.4.2, 11.6, 5.3.2

Description: The link delay asymmetry to correct the link delay.

Default value: 0

NOTE This object is used to verify the topology in P2P.

A.8.3.4.4 Object portDs.profileId

a) Type: ProfileSet

b) Reference: IEC 62439-3:2015

c) Description: The profile in use on this port

Default value: 0

NOTE This object serves to check that all devices are properly configured for deployment or maintenance. This information is not used in the PTP protocol. It does not guarantee that all parameters of the clock have been properly set.

A.8.3.4.5 Object portDs.vlanEnable

Type: Boolean

Reference: IEEE 802.1Q:2014

Description: Enables that PTP messages generated internally to the clock receive an IEEE 802.1Q tag with the values specified in portDs.vlanId and portDs.vlanPrio when sent over this port, with exception of Pdelay_Req and Pdelay_Resp, which operate with link-local addresses. It does not apply to messages in transit generated by other clocks. The bridge rules take precedence over this setting.

Default value: false

NOTE The VLAN settings of the port in its bridge function can override this tagging or remove it. A bridge port cannot send a VID of 0 except if its VLAN setting is disabled.

A.8.3.4.6 Object portDs.vlanId

Type: Unsigned32 (0 to 4095)

Reference: IEEE 802.1Q:2014

Description: The IEEE 802.1Q VLAN Identity to be used by the PTP stack internal to a bridge if portDs.vlanEnable is true.

Default value: 0

NOTE A bridge port would not send a message with a VID of 0 if the port VLAN setting indicates that VLANs are enabled.

A.8.3.4.7 Object portDs.vlanPrio

Type: Unsigned32 (0 to 7)

Reference: IEEE 802.1Q:2104

Definition: The IEEE 802.1Q priority (0 to 7) to be used by PTP messages originating from the PTP stack of this clock if portDs.vlanEnable is true.

Default value: 4

A.8.3.4.8 Object portDs.twoStepFlag

Type: Boolean

Reference: IEC 61588:2009, 8.2.1.2.1

Description: Enables (if true) 2-step correction at sending of this port (if the port implements it) and indicates that this port uses 2-step correction.

Default value: true

NOTE A BC can use a different correction method on different ports, according to the supported profile.

A.8.3.4.9 Object portDs.peerIdentity

Type: OCTET STRING (SIZE (6))

Reference: ISO/IEC/IEEE 8802-3:2014

Description: The Medium Access Control address of the peer clock in a peer-to-peer delay measurement.

A.8.3.4.10 Object portDs.prpAttachment

Type: PortAttachmentType

Reference: IEC 62439-3:2015

Description: The type of redundant attachment of this port.

Default value: 0

A.8.3.4.11 Object portDs.prpPairedPort

Type: Unsigned16 [0 to 4095]

Description: The portNumber of the other port to which this port is attached, if the port is paired for the purpose of double attachment; a value greater than 4 096 indicates the port is not paired.

Default value: 65535

A.8.3.4.12 Object portDs.errorCounter

Type: Counter32

Description: The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that are normally received given the port state.

Default value: 0

A.8.3.4.13 Object portDs.peerDelayLim

Type: PtpTimeInterval

Description: The value of the peer delay that, if exceeded, will cause an event to be raised.

Default value: 100 ns

A.8.4 Object for transparent clocks

A.8.4.1 Additional or modified objects

The members additional or modified with respect to IEC 61588 are as follows:

- transparentClockDefaultDS.profileSet;
- transparentClockDefaultDS.timeInaccuracy
- transparentClockPortDs.portEnabled;
- transparentClockPortDs.dlyAssymetry;transparentClockPortDS.twoStepFlag;
- transparentClockPortDs.peerIdentity;
- transparentClockPortDs.prpAttachment;
- transparentClockPortDs.prpPairedPort;transparentClockPortDs.errorCounter;
- transparentClockPortDs.delayMax.

A.8.4.2 Dataset TransparentClockDefaultDS

A.8.4.2.1 Object transparentClockDefaultDS.profileSet

This object has the same definition as in A.8.3.2.2.

A.8.4.2.2 Object transparentClockDefaultDS.timeInaccuracy

Type: TimeInterval

Units: scaledNanoseconds – IEC 61588:2009, 5.3.2

Reference: IEC 62439-3:2015.

Description: The absolute value of the rated worst-case time error introduced between any two ports.

Default value: 50 ns

NOTE This object serves to check that a TC has a sufficient quality for the given application. This value is not measured at run-time, it is a design parameter that also appears in the MIB. This value is not transmitted using in the PTP protocol.

A.8.4.3 Dataset transparentClockPortDs

A.8.4.3.1 Object transparentClockPortDs.portEnabled

Type: Boolean

Reference: IEC 61588:2009, 8.2.5.3.1

Description: Enables the port for PTP traffic and indicates the enabled status of the port.

Default value: true

A.8.4.3.2 Object transparentClockPortDs.dlyAsymmetry

This object has the same definition as in A.8.3.4.2.

A.8.4.3.3 Object transparentClockPortDs.twoStepFlag

Type: Boolean

Reference: IEC 61588:2009, 8.2.1.2.1

Description: Enables (if true) 2-step correction at sending of this port (if the port implements it) and indicates that this port uses 2-step correction.

Default value: true

NOTE A BC can use a different correction method on different ports, according to the supported profile.

A.8.4.3.4 Object transparentClockPortDs.peerIdentity

Type: OCTET STRING (SIZE (6))

Reference: ISO/IEC/IEEE 8802-3:2014

Description: The Medium Access Control address of the peer clock in a peer-to-peer delay measurement.

A.8.4.3.5 Object transparentClockPortDs.prpAttachment

Type: PortAttachmentType

Reference: IEC 62439-3:2015

Description: The type of redundant attachment of this port.

Default value: 0

A.8.4.3.6 Object transparentClockPortDs.prpPairedPort

Type: Unsigned16 [0 to 4095]

Description: The portNumber of the other port to which this port is attached, if the port is paired for the purpose of double attachment, a value greater than 4 096 indicates the port is not paired.

Default value: 65 535

A.8.4.3.7 Object transparentClockPortDs.errorCounter

Type: Counter32

Description: The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that are normally received given the port state.

Default value: 0

A.8.4.3.8 Object transparentClockPortDs.peerDelayLim

Type: PtpTimeInterval

Description: The value of the peer delay that, if exceeded, will cause an event to be raised.

Default value: 100 ns

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

Annex B
(informative)

Deterministic medium access in HSR

~~This is a placeholder for further development.~~

~~HSR does not specify the traffic control that has to be used for deterministic, real-time operation.~~

~~However, it is a recommended practice to buffer the hard real time, high priority frame, wait until the clock reaches a certain time, the same in all devices with the same period, and let all these nodes send this traffic at that time, in order to leave sufficient contiguous free space for the non-real time traffic.~~

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

Annex B (normative)

PTP profile for Power Utility Automation – Redundant clock attachment

B.1 Application domain

For Power Utility Automation, IEC/IEEE 61850-9-3 specifies a Precision Time Protocol profile using a layer 2 transport according to IEC 61588:2009, Annex F with a peer-to-peer delay measurement based on IEC 61588:2009, Clause J.4.

Annex B defines the mandatory support of redundant clock attachment as an extension of IEC/IEEE 61850-9-3:—, Clause 9.

B.2 PTP profile specification

The profile identifier of a singly or doubly attached clock according to IEC/IEEE 61850-9-3 and Annex A shall be:

profileName: IEC/IEEE 61850-9-3 "Precision Time Protocol profile for Power Utility Automation"

profileVersion: 1.0

profileIdentifier: 00-0C-CD-00-01-xy

Where the first nibble of the sixth octet is a bitfield:

- x = 0 for a singly-attached clock
- x = 1 for PRP redundancy
- x = 2 for HSR redundancy
- x = 3 for both PRP and HSR (configurable) redundancy
- y = 0 (minor revision)

organizationName: IEC Technical Committee 57 Working Group 10

sourceIdentification: A copy of IEC/IEEE 61850-9-3 can be obtained from the International Electrotechnical Commission (<http://www.iec.ch>)

B.3 Redundant clock attachment

Clocks shall implement the doubly attached clock model for at least one of:

- a) PRP (A.4.2) or
- b) HSR (A.5) or
- c) HSR-PRP RedBoxes (A 6)

The manufacturer shall list the supported method in the PICS, as stated in IEC/IEEE 61850-9-3.

Annex C (normative)

PTP profiles for high-availability automation networks

C.1 Application domain

For high-availability industrial automation networks, two precision time clock profiles are specified which share most specifications:

- L3E2E, transported over Layer 3 using end-to-end delay measurement (E2E), based on IEC 61588:2009, Clause J.3 and
- L2P2P, transported over Layer 2 using peer-to-peer delay measurement (P2P), based on IEC 61588:2009, Clause J.4.

An objective is that clock synchronization operates without special configuration if all network elements comply with the same profile.

Interoperability between L3E2E and L2P2P clocks is not an objective, but coupling them through BCs is possible. Clocks of both profiles can coexist on the same network.

These profiles allow mixing singly attached clocks according to the IEC 61588:2009 default profiles and doubly-attached clocks according to Annex A in the same network.

C.2 PTP profile specification

The profile specification of a singly or doubly attached clock according to Annex C shall be according to IEC 61588:2009, 19.3.3:

```

profileName:      IEC 62439-3 "Precision Time Protocol profile for High-availability
                  Automation Networks"

profileVersion:    1.0

profileIdentifier:  00-15-4E-00-01-xy
                  Where the first nibble of the sixth octet is a bit field:
                  x = 1 for PRP redundancy
                  x = 2 for HSR redundancy
                  x = 4 for L3E2E
                  x = 8 for L2P2P
                  y = 0 (minor version)
                  EXAMPLE: 00-15-4E-00-01-50 supports L3E2E and PRP.

organizationName: IEC Technical Committee SC65C Working Group 15

sourceIdentification: A copy of IEC 62439-3 can be obtained from the International
                  Electrotechnical Commission (http://www.iec.ch)
  
```

NOTE The OUI defined in IEC 61588:2009, Clause J.4 does not apply when redundancy is used.

C.3 Clock types

This specification distinguishes clocks by their capabilities:

- Ordinary Clocks (OCs) implement one of the following capabilities:
 - Slave-only (defaultDS.slaveOnly = true)
no port can be in the MASTER state

- Grandmaster-capable (defaultDS.slaveOnly = False); the port (or the port pair in redundancy) can be in the MASTER state
- Grandmaster-only (defaultDS.slaveOnly = False and clockClass = 6 or 7) no port can be in the SLAVE state.
- Transparent Clocks (TCs) correct and forward PTP messages; ports do not have states in IEC 61588:2009.
- Boundary Clocks (BC) are never slave-only and can have either:
 - exactly one port in the SLAVE state and all other in the MASTER or PASSIVE state;
 - all ports in the MASTER state, in which case the BC is the grandmaster.

NOTE Within a time domain, a BC has a port in the SLAVE state in the upper sub-domain and one or several ports in the MASTER state in the lower sub-domain(s) (see IEC TR 61850-90-4).

C.4 Protocol specification common

All clocks shall support the PTP time scale (based on TAI).

All singly attached clocks shall support the default best master clock algorithm in IEC 61588:2009, 9.3.2, 9.3.3 and 9.3.4.

All doubly attached clocks according to IEC 62439-3 shall support in addition the extension to the best master clock algorithm defined in A.7.3.

All clocks shall support both 1-step and 2-step synchronization on ingress, they may use 1-step or 2-step synchronization on egress.

All clocks shall support at least one of the three management mechanisms:

- 1) the alternate management mechanism using the SNMP MIB specified in Annex E and/or
- 2) the alternate management mechanism using the management objects defined in IEC TR 61850-90-4:2013, 19.3 and 19.4 and/or
- 3) the manufacturer-defined fixed values and/or the manufacturer-specific implementation method to access all configurable values as stated in IEC 61588:2009, 15.1.1.

C.5 Protocol specification for L3E2E automation profile

All clocks shall transport PTP messages using Layer 3 (UDP over IP) according to IEC 61588:2009, Annex D (IPv4) or IEC 61588:2009, Annex E (IPv6)

All clocks shall use only multicast addresses, as specified in IEC 61588:2009, Clause D.3 and IEC 61588:2009, Clause E.3 respectively.

All clocks shall support the end-to-end delay calculation defined in IEC 61588:2009, 10.2.

The “portnumber” of all clocks shall be zero in their four most significant bits.

C.6 Protocol specification for L2P2P automation profile

All clocks shall transmit PTP messages on Ethernet layer 2 according to IEC 61588:2009, Annex F.

All clocks shall use multicast communication using the addresses specified in IEC 61588:2009, Clause F.1.

All clocks shall support the peer-to-peer delay measurement defined in IEC 61588:2009, 10.3 and 11.4.

C.7 Timing requirements

C.7.1 Measurement conditions

Steady state is defined as 30 s after a single master starts to send synchronisation messages and 16 s after a change of master, with no change to the environment temperature.

This requirement applies only to clocks that have been energized for 30 min to accommodate for instance temperature-controlled oscillators.

NOTE Assuming a Gaussian distribution, this corresponds to three sigma ($3\sigma = 99.7\%$) or no more than 3 points outside the specified interval, out of 1 000 total points evaluated.

C.7.2 Network time inaccuracy

The following requirements on the network elements aim at achieving a network time inaccuracy better than $\pm 1\ \mu\text{s}$ after crossing approximately 15 TCs or 3 BCs.

C.7.3 Network elements

All network elements shall be clocks according to this standard, conformant media converters and/or conformant links.

C.7.4 Requirements for grandmasters

C.7.4.1 Grandmaster time inaccuracy

A grandmaster-capable clock shall have a time inaccuracy measured between its applied time reference signal and the produced synchronization messages that is smaller than 250 ns.

NOTE This value corresponds to an IEC 61588:2009 clockAccuracy of 22 hex; IEC 61588 defines clockAccuracy.

In case the grandmaster-capable clock has no time reference signal, IEC 61588:2009, J.4.4.1 shall apply.

C.7.4.2 Grandmaster holdover

A grandmaster shall remain within the time inaccuracy of C.7.4.1 for a holdover time of at least 5 s after losing its time reference signal, given that it was in steady state.

C.7.4.3 Grandmaster clockClass in holdover and recovery

A grandmaster clock shall have an IEC 61588:2009 clockClass of:

- 6 while connected to its time reference signal and in steady state;
- 7 after loss of time reference signal, while in holdover;
- 52 when its time error exceeds the value given in C.7.4.1;
- 187 when its time error exceeds $1\ \mu\text{s}$;
- 6 after recovering the time reference signal and in steady state.

NOTE 1 This modifies IEC 61588:2009, Table 5 for the timing requirements of this profile.

NOTE 2 The timeTraceable bit of the flagField indicates if the grandmaster is currently locked to a recognized standard time source or has a stated uncertainty to a recognized standard time source. The clockClass = 6 indicates that the grandmaster is currently locked to its time reference signal. The clock class 6 appears twice, once before loss of time reference signal, and once after recovery of that signal.

C.7.5 Requirements for TCs

A TC shall introduce less than 50 ns of time inaccuracy, measured between the applied synchronization messages at any ingress port and the produced synchronization messages at any egress port, given it is in steady state.

NOTE This time inaccuracy cumulates the time errors in measuring the residence delay, measuring the peer delay in the ingress port and responding to the peer delay measurement from the downstream clocks in the egress port(s). It does not include time error due to asymmetry in either link, nor introduced by the upstream or downstream peer devices involved in the peer delay measurements. This time inaccuracy can be measured using peer devices with known contribution of the peer delay measurement to time error and using links with known or negligible asymmetry, and subtracting those amounts from the observed TC time error.

C.7.6 Requirements for BCs

C.7.6.1 BC time inaccuracy

A BC shall introduce less than 200 ns of time inaccuracy between the port in the SLAVE state and any port in the MASTER state, provided that it is in steady state.

NOTE This time inaccuracy cumulates the time errors in local clock adjustment, measuring the peer delay in the port in the SLAVE state and responding to the peer delay measurement from the downstream clocks in the port(s) in the MASTER state. It does not include time errors due to asymmetry in either link, nor introduced by the upstream or downstream peer devices involved in the peer delay measurements. This time inaccuracy can be measured using peer devices with known contribution of the peer delay measurement to time error and using links with known or negligible asymmetry, and subtracting those amounts from the observed BC time error.

C.7.6.2 BC as free-running grandmaster

In case the BC has no port in the SLAVE state and no time reference signal, IEC 61588:2009, J.4.4.1 shall apply.

C.7.6.3 BC as master in holdover

A BC shall remain within the time inaccuracy of C.7.6.1 for a holdover time of at least 5 s after losing its time reference signal or PTP synchronization, provided that it was in steady state.

C.7.7 Requirements for media converters

Media converters that support IEC 61588 are considered as TCs or BCs and are subject to the requirements of C.7.5 and C.7.6, respectively.

Media converters (e.g. fibre to copper) introduce a delay, considered in peer-to-peer delay measurement. This delay may suffer from a significant jitter and may be different in both directions. Network engineering should use only media converters with known behaviour to estimate the network time inaccuracy at every slave clock.

Media converters shall present a jitter of less than 50 ns and an asymmetry of less than 25 ns.

C.7.8 Requirements for links

Links present a predictable and nearly constant link propagation delay (about 5 μ s/km for fibre or copper), whose average value is regularly calculated by peer-to-peer delay measurement.

NOTE Radio links are not considered.

Since delay asymmetry is not measurable, network engineering must know it to compensate it.

Links shall present a propagation asymmetry of less than 25 ns or shall have a known propagation asymmetry with an asymmetry variation of less than 25 ns.

C.8 Network engineering

To achieve the required network time inaccuracy, careful network design is required, considering the placement of masters and redundant masters and possible network topology changes because of reconfigurations so as not to exceed the number of allowed TCs and BCs.

The network designer should only select network elements knowing their contribution to network time inaccuracy and dependencies on the operating conditions. The network designer should estimate the network time inaccuracy for all slave clocks.

The network designer should use network elements with stricter specifications in more demanding applications or larger networks.

The total time inaccuracy ε_T of the time signal available at a given clock is computed at engineering time as:

$$\varepsilon_T = \varepsilon_{RF} + \varepsilon_{GM} + N_{TC} \times \varepsilon_{TC} + N_{BC} \times \varepsilon_{BC} + N_{MC} \times \varepsilon_{MC}$$

The network time inaccuracy ε_N of the time signal available at a given clock is computed at engineering time as:

$$\varepsilon_N = \varepsilon_{GM} + N_{TC} \times \varepsilon_{TC} + N_{BC} \times \varepsilon_{BC} + N_{MC} \times \varepsilon_{MC}$$

where:

ε_T = total time inaccuracy;

ε_N = network time inaccuracy;

ε_{RF} = worst-case source time inaccuracy;

ε_{GM} = worst-case grandmaster time inaccuracy (see C.7.4);

ε_{TC} = worst-case TC time inaccuracy (see C.7.5);

ε_{BC} = worst-case BC time inaccuracy (see C.7.6);

ε_{MC} = worst-case media converter time inaccuracy (see C.7.7);

N_{TC} = number of TCs in series on the longest path to this clock;

N_{BC} = number of BCs in the path on the longest path to this clock;

N_{MC} = number of media converters in the path on the longest path to this clock.

The network engineering should consider the time error ε_{RF} in the signal received from the recognized standard time source, e.g. the GNSS signal.

The network commissioner should check the time inaccuracy class of the installed network components and verify the topology (e.g. using the management mechanism, if such is provided) and then perform a calibration for the network time inaccuracy of all slave clocks. For this purpose, it is recommended to equip clocks with a 1 PPS output.

C.9 Default settings

All clocks shall be initially configured with the default settings and shall be able to support the ranges (if specified) of Table C.1.

Table C.1 – PTP attributes for the Industrial Automation profile

PTP attribute	Default value	Range
defaultDS.domainNumber	0	as specified in Table 2 of IEC 61588:2009
portDS.logAnnounceInterval	1 for the L3E2E profile 0 for the L2P2P profile	0 to 4 for the L3E2E profile 0 for the L2PTP profile ^a
portDS.logSyncInterval	0	0 ^a
portDS.logMinPdelay_ReqInterval	0	0 ^a
portDS.announceReceiptTimeout	3	3
defaultDS.priority1	255 for slave-only clocks, 128 otherwise	0 to 255 255 for slave-only clocks
defaultDS.priority2	255 for slave-only clocks, 128 otherwise	0 to 255 255 for slave-only clocks
defaultDS.slaveOnly	False (if this parameter is configurable)	True, False
transparentClockdefaultDS.primaryDomain	0	as specified in Table 2 of IEC 61588:2009
τ (Allan deviation sample period)	1,0 s	1,0 s
For each defined range, manufacturers are free to allow wider ranges, which are out of the scope of this profile and which conformance testing will not consider.		
^a Although a master clock is not required to support any other value, it still can have a jitter as IEC 61588:2009, 9.5.9.2 specifies. A slave clock should have a capture range accordingly.		

C.10 Redundant clock handling

Clause C.10 does not apply to clocks with only one port.

These profiles differ from IEC 61588:2009 when clocks are doubly attached according to Annex A as follows:

- 1) Introduction of ordinary clocks with more than one port;
- 2) Extension of the BMCA to allow a decision based on an application-specific clock quality when the same grandmaster appears over two different paths;
- 3) Extension of the port state machine for doubly attached clock to allow checking of unused paths;
- 4) Time-outs for supervision of all message exchanges;
- 5) Extension of the datasets to include redundancy management.

When redundancy is used, it shall be implemented as follows:

Clocks shall implement the doubly attached clock model for at least one of:

- a) PRP (A.4.2) or
- b) HSR (A.5) or
- c) HSR-PRP RedBoxes (A.6)

The manufacturer shall list the supported method in the PICS.

C.11 Protocol Implementation Conformance Statement (PICS)

C.11.1 Conventions

The conventions are based on ISO/IEC 9646-7.

The "Base" column reflects the definitions and specifications in the base standard (i.e. Clauses C.3 to C.10 of this document in this case). Each entry in this column is chosen from the following list:

- **mandatory (m)**: the base standard mandates this capability and it is implemented;
- **optional (o)**: the base standard leaves this capability optional, but it is implemented.

The "Condition" column reflects the condition(s) for the capability. Each entry in this column is chosen from the following list:

- **conditionally supported (c)**: this capability is supported under the conditions specified in the corresponding index in the PICS Table;
- **unconditionally supported (–)**: this capability is not constrained.

The "Support" column is to be filled by the manufacturer.

C.11.2 PICS

The manufacturer shall list the capabilities and options supported by its clock in the PICS according to Table C.2.

The "Base" of this PICS are Clauses C.3 to C.10 of this document.

Table C.2 – PICS for clocks

PICS proforma reference	Capability	Value range	Base	Condition ^a	Support (see Note)
CLOCK_TYPE_OC	clock is OC according to this base	True, False	m	c.1	<i>True</i>
CLOCK_TYPE_TC	clock is TC according to this base	True, False	m	c.1	
CLOCK_TYPE_BC	clock is BC according to this base	True, False	m	c.1	
NR_PORTS	number of clock ports (total)	integer > 0	m	-	2
PORTS_STEP	1: all ports support 1-step on egress 2: all ports support 2-step on egress 3: all ports support both 1-step and 2.	1 to 3	m	-	{1}
SLAVE_ONLY	all ports of the clock are slave-only	True, False	m	c.2	<i>True</i>
TIME_TRACEABLE	connectable to a time reference outside of PTP (e.g. GPS)	True, False	m	c.3	
FREQ_TRACEABLE	connectable to a frequency reference outside of PTP (e.g. GPS)	True, False	m	c.3	
DAC	doubly attached OC	True, False	o	-	<i>True</i>
PORTS_PAISED	paired clock ports for redundancy (e.g. {3-4})	Identifier pair	o	c.4	{0,1}
REDBOX_DATC	Redbox as TC	True, False	o	c.5	
REDBOX_SLTC	Redbox as Stateless TC	True, False	o	c.5	
REDBOX_TWBC	Redbox as three-way BC	True, False	o	c.6	
REDBOX_DABC	Redbox as DAC BC	True, False	o	c.6	
MIB_SNMP	supports MIB of IEC 62439-3, Annex E	True, False	m	c.7	
MIB_61850	supports IEC TR 61850-90-4 Clock Objects	True, False	m	c.7	
MIB_OTHER	clock supports fixed values or a mechanism defined by the manufacturer (If True, this list is appended to this PICS)	True, False	m	c.7	<i>True, void list.</i>
NOTE The values in the “Support” column are an example of a clock that is doubly attached over ports 0 and 1, uses 1-step PdelayReq, but has no master capability and does not support an MIB (there is nothing adjustable). ^a Conditions are: c.1: at least one shall be supported (CLOCK_TYPE_OC and CLOCK_TYPE_TC may be both True) c.2: only if CLOCK_TYPE_OC = True c.3: only if SLAVE_ONLY = False c.4: shall be “m” (>1) if DAC = True c.5: support shall only be declared if CLOCK_TYPE_TC = True and DAC=True c.6: support shall only be declared if CLOCK_TYPE_BC = True and DAC=True c.7: at least one shall be supported					

Annex D (informative)


Precision Time Protocol tutorial for IEC 62439-3

D.1 Objective

This tutorial describes the aspects of the Precision Time Protocol in IEC 61588:2009 that are relevant to IEC 62439-3. D.6.2 describes in addition the conversion between the 1-step and the 2-step correction method, which is currently not specified in IEC 61588:2009, but nevertheless compliant to the standard.

D.2 Precision and accuracy

The quality of a synchronized clock is expressed by its error from an ideal reference time, as Figure D.1 shows.



IEC

Figure D.1 –Precision and accuracy example

The clock accuracy is the mean time error (in Figure D.1: 25 ns).

The precision is the deviation from the mean for a given probability (in Figure D.1: 125 ns at 99,3 % or 50 ns at 68,28 %)

The time inaccuracy of a clock is the deviation from the reference, not from the mean (in Figure D.1: 150 ns at 99,3 %)

Clock accuracy is the maximum guaranteed error from the reference time (in Figure D.1: 200 ns)


Quality criteria for non-synchronized clocks are the stability, expressed for instance in $\mu\text{s/s}$, and the temperature stability ($\mu\text{s/s/}^{\circ}\text{C}$).

A clock that is only synchronized to a reference frequency is called a syntonized clock.

NOTE Similar considerations apply to frequency, which is not considered here.

D.3 PTP clock types

The basic principle of IEC 61588:2009 is shown in Figure D.2.



IEC

Figure D.2 – Precision Time Protocol principle

A grandmaster clock (GMC) distributes the exact time in a Sync message multicast to all (slave) ordinary clocks (OC). This Sync message suffers different delays on the path between the grandmaster and the slaves, especially a residence delay when traversing the bridges and a link delay on the medium.

The bridges operate here as TCs capable of calculating the residence delay and correcting the time for them. To this effect, each TC (and each master and each slave) has hardware capable of time-stamping the PTP messages that transit through them with nanosecond precision.

More generally, PTP distinguishes the following clock types (see Figure D.3):

- Ordinary clocks (OC) have only one port (singly-attached clock), that operates either in the slave or in the master state. An ordinary clock can be configured as master-capable or

slave-only. When the OC is the best clock of its subdomain, it becomes the Grandmaster Clock (GMC); its port is in the MASTER state and sends Sync messages to synchronize all slave clocks. Otherwise, its port is in the SLAVE (or PASSIVE) state and it receives synchronization messages.

- Transparent clocks (TC) have at least two ports which are stateless, and acts as a bridge for all messages. A TC forwards the Sync messages it receives over one port to all its other ports and corrects the time indicated in the Sync messages.
- Boundary clocks (BC) have at least two ports, that are usually in different states, exactly one port being in the SLAVE state and synchronized by a master clock while the other port(s) is (are) in the MASTER state and synchronize(s) a subdomain. When the BC receives a Sync message over the port in the SLAVE state, it does not forward it to the ports in the MASTER state; instead, it sends its own PTP messages, with a period that may be different from port to port. If the BC is the best clock of all subdomains, all its ports are in the MASTER state and it becomes the Grandmaster clock. Otherwise, it becomes master in each subdomain where it is the best master. It may be degraded to SLAVE or PASSIVE by a better master in its subdomain.
- Hybrid clocks (HC) combine a transparent clock (TC) and an ordinary clock (OC).

grandmaster:

top level clock of the time domain
master of the top subdomain,
defines **grandmaster identity**

master port

sends Sync & Announce: sets
"sourcePortIdentity" for this
subdomain

Transparent Clock forwards
and corrects PTP messages

Boundary Clock has a **slave port**
in the upper subdomain and
a **master (or passive)** port in
the lower subdomain

if this link is established, the two subdomains merge and
the **master port** of one boundary clock becomes either
slave or **passive**, according to the BMCA

designated master

(back-up clock of domain)
first to become **master**
(after BMCA) if
the current master stops
sending Announce

slave (or passive) port

ordinary clock can take
the role of **master** or **slave**

Sync & Announce carry
the MAC address of the
master,
(not that of the
grandmaster)
and the
"sourcePortIdentity"
of the **master**

hybrid clocks combine
a TC and an OC,
HC have two or more ports

Key:

The text in red expresses the momentary role of a clock

The text in blue expresses the current state of a port according to its state machine.

The text in **black** is the configured or factory type of clock

Figure D.3 – Precision Time Protocol elements

D.4 PTP main options

IEC 61588:2009 distinguishes three basic options:

- layer 2 versus layer 3 communication (see D.5)
- 1-step versus 2-step clock correction (see D.5).
- end-to-end (see D.7) versus peer-to-peer link delay measurement (see D.8)

D.5 Layer 2 and layer 3 communication

Layer 2 PTP communication transports PTP message on layer 2, relying on MAC addresses. This communication ends at a router. Layer 2 PTP defines for Ethernet at which point in a frame the time stamps is taken. Layer 2 PTP is the preferred approach for automation networks based on Ethernet Layer 2.

Layer 3 PTP communication transports PTP messages over layer 3, relying on IPv4 or IPv6 addresses (mapped to specific Layer 2 addresses). Layer 3 PTP can operate over a variety of media. The precise time-stamping relies on an unspecified physical layer support. Layer 3 PTP is the preferred approach for automation networks based on Layer3.

D.6 1-step and 2-step correction

D.6.1 Time correction in TCs

The master clock (MC) multicasts periodically (e.g. every 1s) a Sync message containing the reference time. This Sync message transits through the network over links and TCs.

The Sync message suffers a path delay consisting of the link delays λ between clocks and of the residence delays ρ in the TCs, before it reaches the slave clocks. To correct these delays, the network devices time-stamp precisely the Sync messages as they ingress and egress the devices (in the physical interface, PHY) as shown in Figure D.4.




Figure D.4 – Delays and time-stamping logic in TCs

The TCs calculate a correction including the residence time of the Sync message. There are two ways to transmit this correction: 1-step correction and 2-step correction, which are illustrated in Figure D.5.

- 1) 1-step correction: the TC enters the value of the correction in the Sync message while it is being sent. This needs a hardware support that can modify messages on-the-fly.
- 2) 2-step correction: the TC sends the value of the correction in a subsequent Follow_Up message. The hardware support is limited to the precise time-stamping of the ingress and egress frames.




Figure D.5 – Correction of the Sync message by 1-step and 2-step (peer-to-peer)

Both methods can be used within a network (see D.6.2).

D.6.2 2-step to 1-step translation

D.6.2.1 Choice of methods

IEC 61588:2009 allows to mix clocks with 1-step correction and 2-step correction within the same subdomain.

The engineer is faced with two correction methods, depending on the ability of its devices. Most current commercial clocks use 2-step correction to simplify implementation, but with increasing silicon support 1-step is gaining ground.

In HSR rings, 1-step is preferred since hardware support is already provided for cut-through and handling Follow_Ups would only be overhead.

Clocks supporting the two correction methods can be chained in the same network if they accept both 1-step and 2-step correction. An ingress port is supposed to accept both 1-step or 2-step corrected messages; it does not need special hardware for this.

A TC forwards the message with 1-step or 2-step correction, possibly modifying the 2-step flag in the message header. A BC treats the clock domains independently.

Although IEC 61588:2009 lists the ability of performing 1-step or 2-step as a default clock property, it is in reality a port property, but usually the same hardware will be used on different ports, so IEC 61588:2009 assumes the same property for all ports.

D.6.2.2 Translation from 1-step to 2-step and vice-versa

A 2-step clock is always able to receive 1-step Sync messages. A 2-step clock that receives a 1-step Sync needs no special translation except change the mode to 2-step and generate the Follow_Up. In one method, the Sync carries the received correction field, while the Follow_Up carries the sum of the residence delay and of the peer delay. Alternatively, the Sync carries the sum of the received correction field and of the peer delay, while the Follow_Up carries only the residence delay. The next clock will take the sum of the correction fields in the Sync and in the Follow_Up messages for correction, as IEC 61588:2009 foresees.

NOTE A 1-step TC modifies the correction field of the Sync message, while a 2-step TC modifies the correction field of the Follow_Up message. Therefore it is important that for any 1-step translation or synchronization the sum of the Sync and Follow_Up correction field is used.

D.6.2.3 Translation from 2-step to 1-step

NOTE 1 IEC 61588:2009 only foresees a 1-step to 2-step translation. The method described in D.6.2.3 is considered for a next revision of IEC 61588, it is fully backward compatible.

Figure D.6 shows the translation from 2-step to 1-step in a chain of 1-step and 2-step TCs between an ordinary clock (slave) and the grandmaster clock (master).




Figure D.6 – Translation from 2-step to 1-step in TCs

A 1-step clock that receives a 2-step Sync awaits the Follow_Up before generating the 1-step Sync. The 1-step Sync at the egress port then carries the sum of the correction fields of the 2-step Sync, of the Follow_Up, the residence delay of the TC and the estimated delay to the upstream peer, as Figure D.7 shows.

NOTE 2 Without synchronized clocks the error can be large. A 500 ms delay between Sync and Follow_Up measured by a 50 ppm local oscillator would cause an error of 25 μ s.




Figure D.7 – Translation from 2-step to 1-step – message view

D.7 End-To-End link delay measurement

D.7.1 General method

In all methods, the residence time of a Sync message in the TCs is measured using the local oscillator, the difference is in the method for measuring the link delays.

In end-to-end link delay measurement, the TCs correct the Sync or Follow_Up messages only for the residence time, but not for the link delays. The calculation of the link delay involves a call-and-reply mechanism using the delay request (Delay_Req) and delay response (Delay_Resp) messages, which yields the sum of all link delays along the path that Sync takes, assuming that link delays are the same in both directions.

To minimize the time error due to the clock drift, the time difference between Sync and Delay_Req is kept short. In practice, Delay_Req is sent in response to a Sync and is evaluated for the next Sync. Path delays do not change rapidly, so this value can be averaged over several periods.

The handling of these messages is different in 1-step and 2-step correction.

D.7.2 End-to-End link delay measurement with 1-step clock correction

Each slave calculates the sum of the link delays on the path to the master by sending a Delay_Req (delay request) message to the master, which responds with a Delay_Resp (delay response) message containing the time of reception. Both Delay_Req and Delay_Resp cross all TCs in the network. The slave awaits the Sync to compute the correction.

The master sends the Sync at time t_1 and the TCs correct the received value of t_1 by their residence time ρ . The slave receive the Sync at time t_2 with the sum of the residence time correction ρ_{ms} .

To calculate the sum of the end-to-end link delays, the slave sends at time t_3 a Delay_Req to the master. The master receives the Delay_Req at time t_4 with the sum of the residence time corrections on the path and responds with a Delay_Resp that contains the received sum of the correction and the timestamp t_4 , as Figure D.8 shows.




Figure D.8 – End-to-end link delay measurement with 1-step clock correction

After this exchange, the slave keeps t_3 , t_4 and the sum of the residence times, $\Sigma \rho_{sm}$.

TCs with 1-step correct the timestamp in the Delay_Req for their residence delay. They do not modify the Delay_Resp.

This mode therefore implies that the Delay_Req and Sync take the same path, which requires some engineering in a layer 3 network. The path taken by Delay_Resp is less critical.

D.7.3 End-to-End link delay measurement with 2-step clock correction

TCs with 2-step correction do not modify the Delay_Req; they just record the exit timestamp of the Delay_Req. When the Delay_Resp comes back, they insert this correction when they forward the Delay_Resp in the opposite direction, as Figure D.9 shows. The Delay_Resp collects all residence delays.

This mode therefore implies that the TCs memorize the communicating partners and that Sync, Delay_Req and Delay_Resp take the same path, which requires some engineering in a layer 3 network.




Figure D.9 – End-to-end delay measurement with 2-step clock correction

D.7.4 End-to-End link delay calculation by Delay_Req/Delay_Resp

In both 1-step and 2-step correction, the slave deduces the sum of the two link delays from the correction field $\Sigma\delta_{sm}$ and $\Sigma\rho_{ms}$ previously received in Delay_Resp and Sync:

$$\Sigma\lambda = \frac{(t_1 - t_4) - (t_2 - t_3) - \Sigma\rho_{ms} - \Sigma\rho_{sm}}{2}$$

Again, this scheme assumes symmetry of the path delays (the residence times can be widely different).

D.8 Peer-to-Peer link delay calculation


D.8.1 Peer-to-Peer link delay calculation with 1-step correction

Unlike end-to-end which measures the sum of the link delays, each peer-to-peer TC computes the link delay to all its neighbours by sending a time-stamped Pdelay_Req (peer delay request) message on its links, to which the peer responds with a time-stamped Pdelay_Resp (peer delay response) message. This delay is summed with the residence delay in the correction of the Sync or Follow_Up.

Figure D.10 shows a PTP clock distribution peer-to-peer link delay measurement and with 1-step correction.

To calculate the link delays, each clock sends peer delay request messages (Pdelay_Req) to all its neighbours, even in the direction opposite to the current GMC since a redundant GMC may exist. The neighbour responds with a Pdelay_Resp message and the link delay over each link is calculated as:

$$\lambda = \frac{(t_4 - t_1) - (t_3 - t_2)}{2}$$



IEC

Figure D.10 – Peer-to-peer link delay measurement with 1-step clock correction

In 1-step clocks, TC_j adds the received correction from clock k to its own correction ($\lambda_{kj} + \rho_j$) and modifies the Sync message's correction field on-the-fly, while the frame is being transmitted over the egress port, since the residence time is only precisely known when transmission started.

This 1-step correction requires a dedicated hardware to modify the correction field (and the checksum) of the outgoing Sync frame while the header of the message has already been transmitted.


The slave clock corrects the received time t_4 by the value of the correction field, which contains the sum of the residence times and the sum of the estimated link delays.

D.8.2 Peer-to-Peer link delay calculation with 2-step correction

To avoid a dedicated hardware for modifying the Sync message on-the-fly, a 2-steps clock records the precise time at which the Sync was received, forwards the Sync message without modification (but announcing a Follow_Up message) and records the precise time at which the Sync was sent. Dedicated PHY chips time-stamp messages within a few nanoseconds, so sending the Sync message is not time-critical, it can be done by a CPU.

The 2-step TCs calculates the correction field as the sum of the ingress Sync correction, the residence delay and the (previously measured) peer delay and sends a Follow_Up message as Figure D.11 shows. The same scheme applies to the calculation of the peer delay.

The receiver assumes that Sync and Follow_Up messages are paired; they have the same PTP sequence number.



IEC

Figure D.11 – Peer-to-peer link delay measurement with 2-step clock correction

A 2-step Pdelay_Resp/Pdelay_Resp_Follow_Up improves over 1-step Pdelay_Resp since it transmits the values of t_2 and t_3 individually and therefore allows to discover errors in the syntonization, as described in IEC 61588:2009, 6.5.5 and specified in IEC 61588:2009, 11.4.3, (c) 8) option (i).

A combination of 2-step with 1-step is in principle possible, since the 2-step flag applies to the messages individually, but is not currently specified. In view of a possible revision of IEC 61588:2009, the receiver should watch the 2-step flag of the received Pdelay_Resp or Sync messages to know if a Pdelay_Resp_Follow_Up or Follow_Up is expected. Also, the revision of IEC 61588:2009 is expected to define a time-out for a Follow_Up to be received, in case of unreliable links.

Annex E (normative)

Management Information base for singly and doubly attached clocks

NOTE When copying the MIB from a PDF document, all headers and footers are removed and characters are converted to the character format of SNMP. Although care has been taken not to use illegal characters, some can appear in the final document due to the editing process.

```

IEC62439-8-MIB DEFINITIONS ::= BEGIN
IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE, Unsigned32, Integer32, Counter32
FROM SNMPv2-SMI          -- [RFC2578]
    TEXTUAL-CONVENTION, TruthValue FROM SNMPv2-TC  -- [RFC2579]
    MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP FROM SNMPv2-CONF --
[RFC2580]
    iec62439 FROM IEC-62439-3-MIB -- root
;
ptp MODULE-IDENTITY
    LAST-UPDATED      "201504081548Z"-- 2015-04-08: 15:48:55
    ORGANIZATION      "IEC/SC 65C WG15"
    CONTACT-INFO      "International Electrotechnical Commission,
Central Office, 3, rue de Varambe 22, CH-1211 Geneva 20, Switzerland"
    DESCRIPTION
        "The Management Information Base module of the IEC62439 Redundant Time
Distribution protocol."
    REVISION           "201504081548Z"-- 2015-04-08: 15:48:55
    DESCRIPTION
        "Original version to be published as part of the CDV"
        ::= { iso(1) std(0) 62439 }
    OBJECT IDENTIFIER ::= { iec62439.ptp(7) }

=====
-- Textual Conventions
=====

ClockIdentity ::= TEXTUAL-CONVENTION
    DISPLAY-HINT    "lx:"
    STATUS          current
    DESCRIPTION
        "Represents an IEEE 802 MAC address represented in the
canonical order defined by IEEE 802.1a, EUI-64. EUI-48
converts to EUI-64 as specified by IEEE. The conversion
assigns values 255 and 254 to octets 3 and 4 respectively,
where octet 0 is the most significant and octet 7 the least.
For example, EUI-48 of AC:DE:48:23:45:67 would extend to
AC:DE:48:FF:FE:23:45:67. In IEC 62439-3 the octets 4 and 5 are used to transport the
GMshortID"
    REFERENCE       "IEC 61588:2009 5.3.4 and 7.5.2.2"
    SYNTAX          OCTET STRING (SIZE (8))

PtpTimeInterval ::= TEXTUAL-CONVENTION
    DISPLAY-HINT    "2x:"
    STATUS          current
    DESCRIPTION
        "TimeInterval Value from IEEE Std 1588-2008 5.3.2 - 2**48 / 10**9 s"
    REFERENCE       "IEC 61588:2009 5.3.2"
    SYNTAX          OCTET STRING (SIZE (8))

ClockClass ::= TEXTUAL-CONVENTION
    STATUS          current
    DESCRIPTION
        "Clock Class Value from IEEE Std 1588-2008 7.6.2.4,
with the following interpretation placed on the value:

```

6: A clock that is synchronized to a primary reference time source,
 7: A clock that has previously been designated as clockClass 6 but that has lost the ability to synchronize to a primary reference time source and is in holdover mode and within holdover specifications,
 13: A clock that is synchronized to an application-specific source of time,
 14: A clock that has previously been designated as clockClass 13 but that has lost the ability to synchronize to an application-specific source of time and is in holdover mode and within holdover specifications,
 52: Degradation alternative A for a clock of clockClass 7 that is not within holdover specification,
 58: Degradation alternative A for a clock of clockClass 14 that is not within holdover specification,
 68..122: For use by alternate PTP profiles (68..122),
 133..170: For use by alternate PTP profiles (133..170),
 187: Degradation alternative B for a clock of clockClass 7 that is not within holdover specification,
 193: Degradation alternative B for a clock of clockClass 14 that is not within holdover specification,
 216..232: For use by alternate PTP profiles,
 248: Default none of the other clockClass definitions apply,
 255: A slave-only clock(255)."

REFERENCE "IEC 61588:2009 7.6.2.4"

SYNTAX INTEGER {
 primarySync(6),
 primarySyncLost(7),
 applicationSpecificSync(13),
 applicationSpecificSyncLost(14),
 primarySyncAlternativeA(52),
 applicationSpecificAlternativeA(58),
 primarySyncAlternativeB(187),
 applicationSpecificAlternativeB(193),
 default(248),
 slaveOnlyClock(255)
 }

ClockAccuracy ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Clock Accuracy Value from IEEE Std 1588-2008 8.6.2.3, with the following

interpretation placed on the value:

32: The time is accurate to within 25 ns,
 33: The time is accurate to within 100 ns,
 34: The time is accurate to within 250 ns,
 35: The time is accurate to within 1 us,
 36: The time is accurate to within 2.5 us,
 37: The time is accurate to within 10 us,
 38: The time is accurate to within 25 us,
 39: The time is accurate to within 100 us,
 40: The time is accurate to within 250 us,
 41: The time is accurate to within 1 ms,
 42: The time is accurate to within 2.5 ms,
 43: The time is accurate to within 10 ms,
 44: The time is accurate to within 25 ms,
 45: The time is accurate to within 100 ms,
 46: The time is accurate to within 250 ms,
 47: The time is accurate to within 1 s,
 48: The time is accurate to within 10 s,
 49: The time is accurate to within > 10 s,
 254: Default indicating unknown"

REFERENCE "IEC 61588:2009 7.6.2.5 and Table 6"

SYNTAX INTEGER {
 timeAccurateTo25ns(32),


```

timeAccurateTo100ns(33),
timeAccurateTo250ns(34),
timeAccurateTo1000ns(35),
timeAccurateTo2500ns(36),
timeAccurateTo10000ns(37),
timeAccurateTo25000ns(38),
timeAccurateTo100000ns(39),
timeAccurateTo250000ns(40),
timeAccurateTo1000000ns(41),
timeAccurateTo2500000ns(42),
timeAccurateTo10ms(43),
timeAccurateTo25ms(44),
timeAccurateTo100ms(45),
timeAccurateTo250ms(46),
timeAccurateTo1s(47),
timeAccurateTo10s(48),
timeAccurateToGT10s(49),
timeAccurateToUnknown(254)
}

```

TimeSourceValue ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"The timeSource is an information only attribute indicating the type of source of time used by a grandmaster clock.

The following value are defined:

```

16: Atomic Clock,
32: GPS,
48: Terrestrial Radio,
57: Serial Time Code, e.g. IRIG-B,
64: PTP,
80: NTP,
96: Hand Set,
144: Other,
160: Internal Oscillator"

```

REFERENCE "IEC 61588:2009 7.6.2.6 and Table 7"

```

SYNTAX INTEGER {
    atomicClock(16),
    gps(32),
    terrestrialRadio(48),
    serialTimeCode(57),
    ptp(64),
    ntp(80),
    handSet(96),
    other(144),
    internalOscillator(160)
}

```

Integer48 ::= TEXTUAL-CONVENTION

DISPLAY-HINT "12d"

STATUS current

DESCRIPTION

"48bit Integer value"

REFERENCE "IEEE 802.3"

SYNTAX OCTET STRING (SIZE (6))

DelayMechanism ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Identifies the path delay measurement mechanism in use"

REFERENCE "IEC 61588:2009 8.2.5.4.4"

```

SYNTAX INTEGER {
    e2e(1),
    p2p(2),
    disabled(254)
}

```

PortStateType ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"The state of the port protocol machine according to 1588 "

REFERENCE "IEC 61588:2009 8.2.5.4.4"

SYNTAX INTEGER {
 undefined (0),
 initializing (1),
 faulty (2),
 disabled (3),
 listening (4),
 preMaster (5),
 master (6),
 passiveMaster (7),
 uncalibrated (8),
 slave (9),
 passiveSlave (10)
 }

ProfileSet ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Identifies the profile supported or in use"

REFERENCE "IEC 62439-3"

SYNTAX BITS {
 resProfile0 (0),
 powerUtilityProfile (7),
 l3e2eProfile (8),
 l2p2pProfile (9)
 }

PortAttachmentType ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"port attachment type according to IEC 62439-3"

REFERENCE "IEC 62439-3"

SYNTAX INTEGER {
 unspecified (0),
 ordinaryClockPort (1),
 boundaryClockPort (2),
 transparentClockPort (3),
 boundaryNode (4),
 doublyAttachedOcPort (5),
 doublyAttachedBcPort (6),
 doublyAttachedTcPort (7),
 pairedStatelessTcPort (8),
 other (255)
 }

TlvTxRxIndex ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"index for transmitted or received TLV object"

REFERENCE "IEC 62439-3"

SYNTAX INTEGER {
 tx (1),
 rx (2)
 }

NetworkProtocol ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"PTP transport protocol in use over that port"

REFERENCE "IEC 61588:2009 7.4.1 Table 3"

SYNTAX INTEGER {

```

    reserved (0),
    udpIpv4 (1),
    udpIPv6 (2),
    ieee8023 (3),
    deviceNet (4),
    controlNet (5),
    profinet (6),
    reserved2 (65535)
}

```

LogTimeInterval ::= TEXTUAL-CONVENTION

DISPLAY-HINT "d"

STATUS current

DESCRIPTION

"Time interval expressed as logarithm base 2 of the number of seconds, e.g. 0 = 1s, 2 = 4s, -3 = 125ms, "

REFERENCE "IEC 61588:2009 7.7.2.1"

SYNTAX Integer32 (-128..127)

```

=====
-- Events
=====

```

ptpNotifications OBJECT IDENTIFIER ::= { ptp 0}

eventFaultyState NOTIFICATION-TYPE

OBJECTS { ptpPortDsPortState }

STATUS current

DESCRIPTION

"A clock detects it has entered a faulty state."

::= { ptpNotifications 1 }

eventRedundancyLoss NOTIFICATION-TYPE

OBJECTS { ptpPortDsPortState }

STATUS current

DESCRIPTION

"A clock that was previously redundantly connected detects loss of one channel"

::= { ptpNotifications 2 }

eventPeerDelayExceeded NOTIFICATION-TYPE

OBJECTS { ptpPortDsPeerDelayLim }

STATUS current

DESCRIPTION

"A port detects an abnormal increase of the peer delay, possibly due to link failure"

::= { ptpNotifications 3 }

```

=====
-- Data Sets
=====

```

ptpObjects OBJECT IDENTIFIER ::= { ptp 1}

ptpOcBc OBJECT IDENTIFIER ::= { ptpObjects 1}

ptpDefaultDs OBJECT IDENTIFIER ::= { ptpOcBc 1}

ptpDefaultDsTwoStepFlag OBJECT-TYPE

SYNTAX TruthValue

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"If all ports of the clock apply two-step time correction: true; otherwise false (The port settings may allow individual two-step correction)"

REFERENCE "IEC61588 8.2.1.2.1"

DEFVAL {true}

::= { ptpDefaultDs 1 }

```
-- .1.1.1

ptpDefaultDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The grandMasterIdentity in the case this clock becomes grandmaster
The octets 4 and 5 should carry the value 0xFFFE."
    REFERENCE    "IEC61588 8.2.1.2.2, 7.6.2.1, 7.5.2.2"
    ::= { ptpDefaultDs 2 }
-- .1.1.2

ptpDefaultDsNumberPorts OBJECT-TYPE
    SYNTAX      Unsigned32(1..255)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The number of PTP ports on the device. (1 if OC, >1 for Boundary Clock).
IEC 62439-3 allows 2 for OC that are doubly attached"
    REFERENCE    "IEC61588 8.2.1.2.3, IEC62439-3 A.8.3.2.1"
    ::= { ptpDefaultDs 3 }
-- .1.1.3

ptpDefaultDsClockClass OBJECT-TYPE
    SYNTAX      ClockClass
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The clockClass specification"
    REFERENCE    "IEC61588 8.2.1.3.1.1, 7.6.2.4."
    ::= { ptpDefaultDs 4 }
-- .1.1.4

ptpDefaultDsClockAccuracy OBJECT-TYPE
    SYNTAX      ClockAccuracy
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The clockAccuracy member of the ptpDefaultDs.clockQuality member"
    REFERENCE    "IEC61588 8.2.1.3.1.2, 5.3.7"
    ::= { ptpDefaultDs 5 }
-- .1.1.5

ptpDefaultDsOffsetScaledLogVaria OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The offsetScaledLogVariance member of the ptpDefaultDs.clockQuality member"
    REFERENCE    "IEC61588 8.2.1.3.1.3, 5.3.7 and 7.6.3.5"
    ::= { ptpDefaultDs 6 }
-- .1.1.6

ptpDefaultDsPriority1 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The priority1 attribute of the local clock"
    REFERENCE    "IEC61588 8.2.1.4.1, 7.6.2.2"
    ::= { ptpDefaultDs 7 }
-- .1.1.7

ptpDefaultDsPriority2 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
```

```

        "The priority2 attribute of the local clock"
        REFERENCE      "IEC61588 8.2.1.4.2, 7.6.2.3"
        ::= { ptpDefaultDs 8 }
-- .1.1.8

ptpDefaultDsDomainNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The domain attribute of the local clock, it applies to all ports in a BC"
        REFERENCE      "IEC61588 8.2.1.4.3, 7.1"
        ::= { ptpDefaultDs 9 }
-- .1.1.9

ptpDefaultDsSlaveOnly OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "If the clock is a slave-only clock: true; otherwise false"
        REFERENCE      "IEC61588 8.2.1.4.4, 9.2.2. and 9.2.3"
        ::= { ptpDefaultDs 10 }
-- .1.1.10

ptpDefaultDsProfileSet OBJECT-TYPE
    SYNTAX      ProfileSet
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The set of profiles which the clock has the ability to support with the
        required performance. Profile activation is a port property"
        REFERENCE      "IEC62439-3 A.8.3.2.2"
        ::= { ptpDefaultDs 11 }
-- .1.1.11

ptpDefaultDsTimeInaccuracy OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The absolute value of the rated worst-case time error with respect to
        absolute time when the clock is connected to its reference signal. If the clock is a
        boundary clock, the reference signal is the clock signal applied to any port, and the
        rated time error the error observed on any master port."
        REFERENCE      "IEC62439-3 A.8.3.2.3"
        ::= { ptpDefaultDs 12 }
-- .1.1.12

ptpDefaultDsOffsetFromMasterLim OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The value of the limit to the offset from master that signals an abnormal
        condition"
        REFERENCE      "IEC62439-3 A.8.3.2.4"
        ::= { ptpDefaultDs 13 }
-- .1.1.13

ptpCurrentDs OBJECT IDENTIFIER ::= { ptpOcBc 2}

ptpCurrentDsStepsRemoved OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION

```

```

        "The number of boundary clocks traversed between the local clock and the
grandmaster clock"
        REFERENCE      "IEC61588 8.2.2.2"
        ::= { ptpCurrentDs 1 }
-- .1.2.1

ptpCurrentDsOffsetFromMaster OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS       "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The implementation-specific, current value of the time difference between a
master and a slave as computed by the slave"
        REFERENCE      "IEC61588 8.2.2.3"
        ::= { ptpCurrentDs 2 }
-- .1.2.2

ptpCurrentDsMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS       "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The implementation-specific, current value of the mean propagation time
between a master and slave clock as computed by the slave"
        REFERENCE      "IEC61588 8.2.2.4"
        ::= { ptpCurrentDs 3 }
-- .1.2.3

ptpParentDs OBJECT IDENTIFIER ::= { ptpOcBc 3}

ptpParentDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The identity of the port on the master that issues the Sync messages"
        REFERENCE      "IEC61588 8.2.3.2"
        ::= { ptpParentDs 1 }
-- .1.3.1

ptpParentDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The number of the port on the master that issues the Sync"
        REFERENCE      "IEC61588 8.2.3.2"
        ::= { ptpParentDs 2 }
-- .1.3.2

ptpParentDsStats OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "If all of the following conditions are satisfied:
- The clock has a port in the SLAVE state.
- The clock has computed statistically valid estimates of the
ptpParentDs.observedParentOffsetScaledLog Variance and the
ptpParentDs.observedParentClockPhaseChangeRate members: true
otherwise false."
        REFERENCE      "IEC61588 8.2.3.3"
        ::= { ptpParentDs 3 }
-- .1.3.3

ptpParentDsObsOfstScdLVar OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only

```

```
STATUS      current
DESCRIPTION
    "If ptpParentDs.parentStats is true, an estimate PTP variance of the parent
    clock as observed by the slave clock. The computation of this value is optional;
    If ptpParentDs.parentStats is false: not valid"
REFERENCE    "IEC61588 8.2.3.4, 7.6.3.5, 7.6.3.3"
::= { ptpParentDs 4 }
-- .1.3.4

ptpParentDsObsPhChgRate OBJECT-TYPE
SYNTAX      Integer32
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The estimate of the phase change rate of the parent clock
    (ptpParentDs.observedParentClockPhaseChangeRate) as observed by the slave clock. If
    the estimate exceeds the capacity of its data type, this value is set to 7FFF FFFF16
    or 8000 000016, as appropriate."
REFERENCE    "IEC61588 8.2.3.5, 7.6.4.4, and 15.5.3.5.1.4"
::= { ptpParentDs 5 }
-- .1.3.5

ptpParentDsGmIdentity OBJECT-TYPE
SYNTAX      ClockIdentity
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The received grandMasterIdentity"
REFERENCE    "IEC61588 8.2.3.6, 7.5.2.2, and 5.3.4"
::= { ptpParentDs 6 }
-- .1.3.6

ptpParentDsGmClass OBJECT-TYPE
SYNTAX      ClockClass
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The clockQuality.clockClass attribute of the grandmaster clock."
REFERENCE    "IEC61588 8.2.3.7, 7.6.2.4"
::= { ptpParentDs 7 }
-- .1.3.8

ptpParentDsGmAccuracy OBJECT-TYPE
SYNTAX      ClockAccuracy
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The clockQuality.clockAccuracy attribute of the grandmaster clock."
REFERENCE    "IEC61588 8.2.3.7, 7.6.2.5"
::= { ptpParentDs 8 }
-- .1.3.9

ptpParentDsGmOfstScdLVar OBJECT-TYPE
SYNTAX      Integer32
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The clockQuality.offsetScaledLogVariance attribute of the grandmaster
    clock."
REFERENCE    "IEC61588 8.2.3.7, 7.6.3, 7.6.2.2"
::= { ptpParentDs 9 }
-- .1.3.10

ptpParentDsGmPriority1 OBJECT-TYPE
SYNTAX      Unsigned32(0..255)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
    "The priority1 attribute of the grandmaster clock."
REFERENCE    "IEC61588 8.2.3.8, 7.6.2.2"
```

```

    ::= { ptpParentDs 10 }
-- .1.3.11

ptpParentDsGmPriority2 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The priority2 attribute of the grandmaster clock."
    REFERENCE    "IEC61588 8.2.3.9, 7.6.2.3"
    ::= { ptpParentDs 11 }
-- .1.3.12

ptpTimePropDs OBJECT IDENTIFIER ::= { ptpOcBc 4}

ptpTimePropDsCurrentUtcOffset OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and if the epoch is the PTP epoch, the
value of defaultDS.currentUtcOffset (offset between TAI and UTC in seconds);
otherwise the value has no meaning"
    REFERENCE    "IEC61588 8.2.4.2, 7.2.3, 9.4"
    ::= { ptpTimePropDs 1 }
-- .1.4.1

ptpTimePropDsCurrentUtcOffsetVld OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and if the defaultDS.currentUtcOffset is
known to be correct: true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.3, 7.2.3, 9.4"
    ::= { ptpTimePropDs 2 }
-- .1.4.2

ptpTimePropDsLeap59 OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and the epoch is the PTP epoch and the last
minute of the current UTC day contains 59 seconds: true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.4, 9.4"
    ::= { ptpTimePropDs 3 }
-- .1.4.3

ptpTimePropDsLeap61 OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and the epoch is the PTP epoch and the last
minute of the current UTC day contains 61 seconds (timePropertiesDS.leap61): true;
otherwise false"
    REFERENCE    "IEC61588 8.2.4.5, 9.4"
    ::= { ptpTimePropDs 4 }
-- .1.4.4

ptpTimePropDsTimeTraceable OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the timescale is PTP and the source is traceable to a primary reference:
true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.6"
    ::= { ptpTimePropDs 5 }

```


-- .1.4.5

```
ptpTimePropDsFrequencyTraceable OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the frequency determining the timescale is traceable to a primary
reference: true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.7"
    ::= { ptpTimePropDs 6 }
-- .1.4.6
```

```
ptpTimePropDsPtpTimescale OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock timescale of the grandmaster clock is PTP: true; otherwise
false"
    REFERENCE    "IEC61588 8.2.4.8, 7.2.1"
    ::= { ptpTimePropDs 7 }
-- .1.4.7
```

```
ptpTimePropDsTimeSource OBJECT-TYPE
    SYNTAX      TimeSourceValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The source of time used by the grandmaster clock given by the enumeration
type"
    REFERENCE    "IEC61588 8.2.4.9"
    ::= { ptpTimePropDs 8 }
-- .1.4.8
```

```
ptpForeignMasterDsTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF PtpForeignMasterDsEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A table of foreign master entries"
    REFERENCE    "IEC61588 9.3.2"
    ::= { ptpOcBc 5 }
-- .1.5
```

```
ptpForeignMasterDsEntry OBJECT-TYPE
    SYNTAX      PtpForeignMasterDsEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A list of foreign master entries"
    REFERENCE    "IEC61588 9.3.2"
    INDEX        {
        ptpForeignMasterDsIndex }
    ::= { ptpForeignMasterDsTable 1 }
-- .1.5.1
```

```
PtpForeignMasterDsEntry ::= SEQUENCE {
    ptpForeignMasterDsIndex      Unsigned32,
    ptpForeignMasterDsPortIdentity ClockIdentity,
    ptpForeignMasterAnnounceMessages Unsigned32
}
```

```
ptpForeignMasterDsIndex OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "Index for ptpForeignMasterDs Table"
```

```

REFERENCE      "IEC61588 9.3.2.4.5"
::= { ptpForeignMasterDsEntry 1 }
-- .1.5.1.1

ptpForeignMasterDsPortIdentity OBJECT-TYPE
SYNTAX      ClockIdentity
MAX-ACCESS   read-only
STATUS       current
DESCRIPTION
    "The sourcePortIdentity field value of an Announce message received from the
    foreign master"
REFERENCE      "IEC61588 9.3.2.4.2"
::= { ptpForeignMasterDsEntry 2 }
-- .1.5.1.2

ptpForeignMasterAnnounceMessages OBJECT-TYPE
SYNTAX      Unsigned32
MAX-ACCESS   read-only
STATUS       current
DESCRIPTION
    "The number of Announce messages from the foreign master indicated by the
    foreignMasterDS.foreignMasterPortIdentity member
    that have been received within a time window of duration FOREIGN_MASTER_TIME_WINDOW"
REFERENCE      "IEC61588 9.3.2.4.3"
::= { ptpForeignMasterDsEntry 3 }
-- .1.5.1.2

ptpPortDsTable OBJECT-TYPE
SYNTAX      SEQUENCE OF PtpPortDsEntry
MAX-ACCESS   not-accessible
STATUS       current
DESCRIPTION
    "A table of port related variables, a value of 1 is used in devices with only
    one port"
::= { ptpOcBc 6 }
-- .1.6

ptpPortDsEntry OBJECT-TYPE
SYNTAX      PtpPortDsEntry
MAX-ACCESS   not-accessible
STATUS       current
DESCRIPTION
    "A list of objects pertaining to a port."
REFERENCE      "IEC61588 8.2.5"
INDEX        {
    ptpPortDsIndex }
::= { ptpPortDsTable 1 }
-- .1.6.1

PtpPortDsEntry ::= SEQUENCE {
    ptpPortDsIndex      Unsigned32,
    ptpPortDsClockIdentity ClockIdentity,
    ptpPortDsPortNumber Unsigned32,
    ptpPortDsPortState   PortStateType,
    ptpPortDsLogMinDelayReqInterval LogTimeInterval,
    ptpPortDsPeerMeanLinkDelay PtpTimeInterval,
    ptpPortDsLogAnnounceInterval LogTimeInterval,
    ptpPortDsAnnounceReceiptTimeout Unsigned32,
    ptpPortDsLogSyncInterval LogTimeInterval,
    ptpPortDsDelayMechanism DelayMechanism,
    ptpPortDsLogMinPdelayReqInterval LogTimeInterval,
    ptpPortDsVersionNumber Unsigned32,
    ptpPortDsNetworkProtocol NetworkProtocol,
    ptpPortDsPortEnabled TruthValue,
    ptpPortDsDlyAsymmetry PtpTimeInterval,
    ptpPortDsProfileId ProfileSet,
    ptpPortDsVlanEnable TruthValue,
    ptpPortDsVlanId Unsigned32,
    ptpPortDsVlanPrio Unsigned32,

```

```

        ptpPortDsTwoStepFlag      TruthValue,
        ptpPortDsPeerIdentity      ClockIdentity,
        ptpPortDsPrpAttachement    PortAttachmentType,
        ptpPortDsPrpPairedPort      Unsigned32,
        ptpPortDsErrorCounter        Counter32,
        ptpPortDsPeerDelayLim        PtpTimeInterval
    }

ptpPortDsIndex OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "Index for ptpPortDs Table"
    ::= { ptpPortDsEntry 1 }
-- .1.6.1.1

ptpPortDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The portIdentity.clockIdentity attribute of the local port"
    REFERENCE   "IEC61588 8.2.5.2.1, 7.5.2.2.1"
    ::= { ptpPortDsEntry 2 }
-- .1.6.1.2

ptpPortDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The portIdentity.portNumber attribute of the local port"
    REFERENCE   "IEC61588 8.2.5.2.1, 7.5.2.3"
    ::= { ptpPortDsEntry 3 }
-- .1.6.1.3

ptpPortDsPortState OBJECT-TYPE
    SYNTAX      PortStateType
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The current state of the protocol engine associated with this port taken
        from the enumeration in Table 8, extended by the PASSIVE_SLAVE state in doubly
        attached clocks"
    REFERENCE   "IEC61588 8.2.5.3.1, 9.2; IEC62439-3 A.8.3.4.1"
    ::= { ptpPortDsEntry 4 }
-- .1.6.1.4

ptpPortDsLogMinDelayReqInterval OBJECT-TYPE
    SYNTAX      LogTimeInterval
    UNITS        "seconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The logarithm to the base 2 of the minDelayReqInterval; The initialization
        value is implementation-specific consistent with 7.7.2.4."
    REFERENCE   "IEC61588 8.2.5.3.2, 7.7.2.4"
    ::= { ptpPortDsEntry 5 }
-- .1.6.1.5

ptpPortDsPeerMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION

```

"If the value of the ptpPortDs.delayMechanism member is peer-to-peer (P2P), an estimate of the current one-way propagation delay on the link, i.e., <meanPathDelay>, attached to this port computed using the peer delay mechanism. If the value of the ptpPortDs.delayMechanism member is end-to-end (E2E), this value is zero"

```
REFERENCE      "IEC61588 8.2.5.3.3, 11.4"
::= { ptpPortDsEntry 6 }
-- .1.6.1.6
```

ptpPortDsLogAnnounceInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the mean announceInterval"
REFERENCE   "IEC61588 8.2.5.4.1, 7.7.2.2"
::= { ptpPortDsEntry 7 }
-- .1.6.1.7
```

ptpPortDsAnnounceReceiptTimeout OBJECT-TYPE

```
SYNTAX      Unsigned32(0..255)
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The integral multiple of announceInterval"
REFERENCE   "IEC61588 8.2.5.4.2, 7.7.3.1"
::= { ptpPortDsEntry 8 }
-- .1.6.1.8
```

ptpPortDsLogSyncInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the mean SyncInterval for multicast messages"
REFERENCE   "IEC61588 8.2.5.4.3, 7.7.2.3"
::= { ptpPortDsEntry 9 }
-- .1.6.1.9
```

ptpPortDsDelayMechanism OBJECT-TYPE

```
SYNTAX      DelayMechanism
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The propagation delay measuring option used by the port in computing
    <meanLinkDelay>, the data type being enumeration taken from Table 9. The
    initialization value is implementation-specific unless otherwise stated."
REFERENCE   "IEC61588 8.2.5.4.4"
::= { ptpPortDsEntry 10 }
-- .1.6.1.10
```

ptpPortDsLogMinPdelayReqInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the minPdelayReqInterval"
REFERENCE   "IEC61588 8.2.5.4.5, 7.7.2.5"
::= { ptpPortDsEntry 11 }
-- .1.6.1.11
```

ptpPortDsVersionNumber OBJECT-TYPE

```
SYNTAX      Unsigned32(2..15)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
```

```
"The PTP version in use on the port; it is 2 for IEC 61588-2009"
REFERENCE      "IEC61588 8.2.5.4.6"
::= { ptpPortDsEntry 12 }
-- .1.6.1.12

ptpPortDsNetworkProtocol OBJECT-TYPE
    SYNTAX      NetworkProtocol
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The network protocol in use over that port according to Table 3"
    REFERENCE    "IEC61588 7.4.1"
    ::= { ptpPortDsEntry 13 }
-- .1.6.1.13

ptpPortDsPortEnabled OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "Enables the port for PTP traffic and indicates the enabled status of the
port if true"
    REFERENCE    "IEC62439-3 A.8.3.4.2"
    DEFVAL       {1}
    ::= { ptpPortDsEntry 14 }
-- .1.6.1.14

ptpPortDsDlyAsymmetry OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The link delay asymmetry to correct the link delay"
    REFERENCE    "IEC61588 7.4.2, IEC62439-3 A.8.3.4.3"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpPortDsEntry 15 }
-- .1.6.1.15

ptpPortDsProfileId OBJECT-TYPE
    SYNTAX      ProfileSet
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The profiles in use on this port"
    REFERENCE    "IEC62439-3 A.8.3.4.4"
    ::= { ptpPortDsEntry 16 }
-- .1.6.1.16

ptpPortDsVlanEnable OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "Enables that PTP messages generated internally to the clock receive an
802.1Q tag with the values specified in portDs.vlanId and portDs.vlanPrio when sent
over this port, with exception of PdelayReq and PdelayResp, which operate with link-
local addresses. It does not apply to messages in transit generated by other clocks.
The bridge rules take precedence over this setting.
This does not concern PTP messages in transit (even modified), which are handled
according to the bridging rules.
In a boundary clock, this setting applies to the sending over that port only."
    REFERENCE    "IEC62439-3 A.8.3.4.5; IEEE 802.1Q"
    ::= { ptpPortDsEntry 17 }
-- .1.6.1.17

ptpPortDsVlanId OBJECT-TYPE
    SYNTAX      Unsigned32(0..4095)
    MAX-ACCESS   read-write
    STATUS      current
```

```

DESCRIPTION
    "The 802.1Q VLAN Identity to be used by the PTP stack internal to a bridge if
portDs.vlanEnable is true; A bridge port would not send a message with a VID of 0 if
the port VLAN setting indicates that VLANs are enabled"
REFERENCE    "IEC62439-3 A.8.3.4.6; IEEE 802.1Q"
    ::= { ptpPortDsEntry 18 }
-- .1.6.1.18

ptpPortDsVlanPrio OBJECT-TYPE
    SYNTAX      Unsigned32(0..7)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The 802.1Q priority (0 to 7) to be used by PTP messages originating from the
PTP stack of this clock if portDs.vlanEnable is true"
    REFERENCE    "IEC62439-3 A.8.3.4.7; IEEE 802.1Q"
    ::= { ptpPortDsEntry 19 }
-- .1.6.1.19

ptpPortDsTwoStepFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "Enables (if true) 2-step correction at sending of this port (if the port
implements it) and indicates that this port uses 2-step correction"
    REFERENCE    "IEC62439-3 A.8.3.4.8"
    ::= { ptpPortDsEntry 20 }
-- .1.6.1.20

ptpPortDsPeerIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The Medium Access Control address of the peer clock in a peer-to-peer delay
measurement"
    REFERENCE    "IEC62439-3 A.8.3.4.9"
    ::= { ptpPortDsEntry 21 }
-- .1.6.1.21

ptpPortDsPrpAttachement OBJECT-TYPE
    SYNTAX      PortAttachementType
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The type of redundant attachment of this port."
    REFERENCE    "IEC62439-3 A.8.3.4.10"
    ::= { ptpPortDsEntry 22 }
-- .1.6.1.22

ptpPortDsPrpPairedPort OBJECT-TYPE
    SYNTAX      Unsigned32(0..4095)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The portNumber of the other port to which this port is attached, if the port
is paired for the purpose of double attachment; a value greater than 4095 indicates
the port is not paired"
    REFERENCE    "IEC62439-3 A.8.3.4.11"
    ::= { ptpPortDsEntry 23 }
-- .1.6.1.23

ptpPortDsErrorCounter OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION

```

"The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that should be normally received given the port state"

REFERENCE "IEC62439-3 A.8.3.4.12"
 ::= { ptpPortDsEntry 24 }
 -- .1.6.1.24

ptpPortDsPeerDelayLim OBJECT-TYPE
 SYNTAX PtpTimeInterval
 UNITS "scaledNanoseconds"
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The value of the peer delay that, if exceeded, will cause an event to be raised"

REFERENCE "IEC62439-3 A.8.3.4.13"
 DEFVAL {'640000'H}
 ::= { ptpPortDsEntry 25 }
 -- .1.6.1.25

ptpTc OBJECT IDENTIFIER ::= { ptpObjects 2 }

ptpTcDefaultDs OBJECT IDENTIFIER ::= { ptpTc 1 }

ptpTcDefaultDsClkIdentity OBJECT-TYPE
 SYNTAX ClockIdentity
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The clockIdentity attribute of the local clock."

REFERENCE "IEC61588 8.3.2.2.1, 7.6.2.1"
 ::= { ptpTcDefaultDs 1 }
 -- .2.1.2

ptpTcDefaultDsNumberPorts OBJECT-TYPE
 SYNTAX Unsigned32(1..4095)
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The number of PTP ports of the clock (1 if OC), 2 in IEC 62439-3 if Doubly Attached OC and the port is paired"

REFERENCE "IEC61588 8.3.2.2.2"
 ::= { ptpTcDefaultDs 2 }
 -- .2.1.3

ptpTcDefaultDsDelayMech OBJECT-TYPE
 SYNTAX DelayMechanism
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The mechanism used to measure the link delay.

E2E according to Table 9 if the transparent clock uses the end-to-end delay measurement;

P2P according to Table 9 if the transparent clock uses the peer-to-peer delay measurement"

REFERENCE "IEC61588 8.3.2.3.1"
 DEFVAL {2}
 ::= { ptpTcDefaultDs 3 }
 -- .2.1.4

ptpTcDefaultDsPriDomain OBJECT-TYPE
 SYNTAX Unsigned32(0..255)
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The value of transparentClockDefaultDs.primaryDomain of the primary syntonization domain; The initialization value is 0"

REFERENCE "IEC61588 8.3.2.3.2, 10.1"
 ::= { ptpTcDefaultDs 4 }
 -- .2.1.5

```

ptpTcDefaultDsProfileSet OBJECT-TYPE
    SYNTAX      ProfileSet
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The set of profiles which the clock has the ability to support with the
        required performance (Profile activation is a port property)"
    REFERENCE    "IEC62439-3 A.8.4.2.1"
    ::= { ptpTcDefaultDs 5 }
-- .2.1.6

ptpTcDefaultDsTimeInaccuracy OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The absolute value of the rated worst-case time error introduced between any
        two ports"
    REFERENCE    "IEC62439-3 A.8.4.2.2"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpTcDefaultDs 6 }
-- .2.1.7

ptpTcPortDsTable OBJECT-TYPE
    SYNTAX      SEQUENCE OF PtpTcPortDsEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "The table of port dataset related variables. A value of 1 is used in devices
        with only one port."
    ::= { ptpTc 2 }
-- .2.2

ptpTcPortDsEntry OBJECT-TYPE
    SYNTAX      PtpTcPortDsEntry
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION
        "A list of objects pertaining to a port."
    REFERENCE    "IEC61588 8.2.5"
    INDEX        {
        ptpTcPortDsIndex }
    ::= { ptpTcPortDsTable 1 }
-- .2.2.1

PtpTcPortDsEntry ::= SEQUENCE {
    ptpTcPortDsIndex      Unsigned32,
    ptpTcPortDsClockIdentity ClockIdentity,
    ptpTcPortDsPortNumber Unsigned32,
    ptpTcPortDsLogMinPdelayReqInterv Integer32,
    ptpTcPortDsFaultyFlag TruthValue,
    ptpTcPortDsPeerMeanLinkDelay PtpTimeInterval,
    ptpTcPortDsPortEnabled TruthValue,
    ptpTcPortDsDlyAsymmetry PtpTimeInterval,
    ptpTcPortDsTwoStepFlag TruthValue,
    ptpTcPortDsPeerIdentity ClockIdentity,
    ptpTcPortDsPrpAttachement PortAttachmentType,
    ptpTcPortDsPrpPairedPort Unsigned32,
    ptpTcPortDsErrorCounter Counter32,
    ptpTcPortDsPeerDelayLim PtpTimeInterval
}

ptpTcPortDsIndex OBJECT-TYPE
    SYNTAX      Unsigned32 (0..65535)
    MAX-ACCESS   not-accessible
    STATUS       current
    DESCRIPTION

```



```

        "Port index"
        ::= { ptpTcPortDsEntry 1 }
-- .2.2.1.1

ptpTcPortDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The portIdentity.clockIdentity attribute of the local port"
    REFERENCE    "IEC61588 8.3.3.2.1, 7.5.2.2.1"
    ::= { ptpTcPortDsEntry 2 }
-- .2.2.1.2

ptpTcPortDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The portIdentity.portNumber attribute of the local port"
    REFERENCE    "IEC61588 8.3.3.2.1, 7.5.2.3"
    ::= { ptpTcPortDsEntry 3 }
-- .2.2.1.3

ptpTcPortDsLogMinPdelayReqInterv OBJECT-TYPE
    SYNTAX      Integer32(-128..127)
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The value of transparentClockPortDs.logMinPdelayReqInterval as the logarithm
to the base 2 of the
minPdelayReqInterval"
    REFERENCE    "IEC61588 8.3.3.3.1, 7.7.2.5"
    ::= { ptpTcPortDsEntry 4 }
-- .2.2.1.4

ptpTcPortDsFaultyFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "If the transparentClockPortDs.faultyFlag is set: true; otherwise: false
(initialization value)"
    REFERENCE    "IEC61588 8.3.3.3.2"
    ::= { ptpTcPortDsEntry 5 }
-- .2.2.1.5

ptpTcPortDsPeerMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "If the transparentClockDefaultDs.delayMechanism is P2P, the value of
transparentClockPortDs.peerMeanPathDelay as an estimate of the current one-way
propagation
delay on the link attached to this port computed using the peer delay mechanism;
otherwise, if the value of the transparentClockDefaultDs.delayMechanism member is
E2E, the value
is zero. The data type should be TimeInterval."
    REFERENCE    "IEC61588 8.3.3.3.3, 11.4"
    DEFVAL       {'FFFFFFFF'H}
    ::= { ptpTcPortDsEntry 6 }
-- .2.2.1.6

ptpTcPortDsPortEnabled OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION

```

```

        "Enables the port for PTP traffic and indicates the enabled status of the
port if true"
        REFERENCE      "IEC62439-3 A.8.4.3.1"
        DEFVAL          {1}
        ::= { ptpTcPortDsEntry 7 }
-- .2.2.1.7

ptpTcPortDsDlyAsymmetry OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The link delay asymmetry to correct the link delay"
    REFERENCE    "IEC61588 7.4.2, IEC62439-3 A.8.4.3.2"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpTcPortDsEntry 8 }
-- .2.2.1.8

ptpTcPortDsTwoStepFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "Enables (if true) 2-step correction at sending of this port (if the port
implements it) and indicates that this port uses 2-step correction"
    REFERENCE    "IEC62439-3 A.8.4.3.3"
    DEFVAL       {1}
    ::= { ptpTcPortDsEntry 9 }
-- .2.2.1.9

ptpTcPortDsPeerIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The Medium Access Control address of the peer clock in a peer-to-peer delay
measurement"
    REFERENCE    "IEC62439-3 A.8.4.3.4"
    ::= { ptpTcPortDsEntry 10 }
-- .2.2.1.10

ptpTcPortDsPrpAttachement OBJECT-TYPE
    SYNTAX      PortAttachmentType
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The type of redundant attachment of this port"
    REFERENCE    "IEC62439-3 A.8.4.3.5"
    DEFVAL       {3}
    ::= { ptpTcPortDsEntry 11 }
-- .2.2.1.11

ptpTcPortDsPrpPairedPort OBJECT-TYPE
    SYNTAX      Unsigned32(0..4095)
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The portNumber of the other port to which this port is attached, if the port
is paired for the purpose of double attachment; a value greater than 4095 indicates
the port is not paired"
    REFERENCE    "IEC62439-3 A.8.4.3.6"
    DEFVAL       {0}
    ::= { ptpTcPortDsEntry 12 }
-- .2.2.1.12

ptpTcPortDsErrorCounter OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS       current

```

```

DESCRIPTION
    "The error counter that is incremented for each reception error of PTP
    messages (e.g. time-out), only for those messages that should be normally received
    given the port state."
REFERENCE      "IEC62439-3 A.8.4.3.7"
::= { ptpTcPortDsEntry 13 }
-- .2.2.1.13

ptpTcPortDsPeerDelayLim OBJECT-TYPE
SYNTAX      PtpTimeInterval
UNITS       "scaledNanoseconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The value of the peer delay that, if exceeded, will cause an event to be
    raised."
REFERENCE      "IEC62439-3 A.8.4.3.8"
DEFVAL      {'640000'H}
::= { ptpTcPortDsEntry 14 }
-- .2.2.1.14

ptpTLVs OBJECT IDENTIFIER ::= { ptpObjects 3 }

ptpAlternateTimeOffsetIndicator OBJECT IDENTIFIER ::= { ptpTLVs 1 }

ptpAtoiTlvTable OBJECT-TYPE
SYNTAX      SEQUENCE OF PtpAtoiTlvEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "Table of ALTERNATE_TIME_OFFSET_INDICATORS"
REFERENCE      "IEC61588 16.3"
::= { ptpAlternateTimeOffsetIndicator 1 }
-- .3.1.1

ptpAtoiTlvEntry OBJECT-TYPE
SYNTAX      PtpAtoiTlvEntry
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "Table entry"
REFERENCE      "IEC61588 16.3"
INDEX      {
    ptpAtoiTlvIndex,
    ptpAtoiTlvTxRxIndex }
::= { ptpAtoiTlvTable 1 }
-- .3.1.1.1

PtpAtoiTlvEntry ::= SEQUENCE {
    ptpAtoiTlvIndex      Unsigned32,
    ptpAtoiTlvTxRxIndex  TlvTxRxIndex,
    ptpAtoiCurrentOffset  Integer32,
    ptpAtoiJumpSeconds    Integer32,
    ptpAtoiTimeOfNextJump Integer48,
    ptpAtoiDisplayName    OCTET STRING
}

ptpAtoiTlvIndex OBJECT-TYPE
SYNTAX      Unsigned32 (0..65535)
MAX-ACCESS  not-accessible
STATUS      current
DESCRIPTION
    "Table Index for each alternate time offset indicator"
REFERENCE      "IEC61588 16.3"
::= { ptpAtoiTlvEntry 1 }
-- .3.1.1.1.1

ptpAtoiTlvTxRxIndex OBJECT-TYPE
SYNTAX      TlvTxRxIndex

```

```

MAX-ACCESS    not-accessible
STATUS        current
DESCRIPTION
    "Table Index for received or sent objects"
REFERENCE      "IEC61588 16.3"
::= { ptpAtoiTlvEntry 2 }
-- .3.1.1.1.2

ptpAtoiCurrentOffset OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    UNITS        "seconds"
    MAX-ACCESS    read-write
    STATUS        current
    DESCRIPTION
        "The value of currentOffset (offset of the alternate time, in seconds, from
the time of this node). The alternate time is the sum of this value and the time of
the node, which is TAI if the PTP time scale is used."
    REFERENCE      "IEC61588 16.3.3.4"
    ::= { ptpAtoiTlvEntry 3 }
-- .3.1.1.1.3

ptpAtoiJumpSeconds OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    UNITS        "seconds"
    MAX-ACCESS    read-write
    STATUS        current
    DESCRIPTION
        "The value of jumpSeconds (size of the next discontinuity, in seconds, of the
alternate time). A value of zero indicates that no discontinuity is expected. A
positive value indicates that the discontinuity
will cause the currentOffset of the alternate time to increase."
    REFERENCE      "IEC61588 16.3.3.5"
    ::= { ptpAtoiTlvEntry 4 }
-- .3.1.1.1.4

ptpAtoiTimeOfNextJump OBJECT-TYPE
    SYNTAX      Integer48
    UNITS        "seconds"
    MAX-ACCESS    read-write
    STATUS        current
    DESCRIPTION
        "The value of timeOfNextJump (value of the seconds portion of the time of the
transmitting node at the time that the next discontinuity will occur). The
discontinuity occurs at the start of the second indicated
by the value of timeOfNextJump."
    REFERENCE      "IEC61588 16.3.3.6"
    ::= { ptpAtoiTlvEntry 5 }
-- .3.1.1.1.5

ptpAtoiDisplayName OBJECT-TYPE
    SYNTAX      OCTET STRING (SIZE (10))
    MAX-ACCESS    read-write
    STATUS        current
    DESCRIPTION
        "The value of displayName (text name of the alternate timescale in ASCII-8).
It represent the time zone, e.g. CET;
a leading character of * indicates that daylight saving time is in effect."
    REFERENCE      "IEC61588 16.3.3.7"
    ::= { ptpAtoiTlvEntry 6 }
-- .3.1.1.1.6

=====
-- Conformance Information
=====

ptpConformance OBJECT IDENTIFIER ::= { ptp 2}
ptpGroups OBJECT IDENTIFIER ::= { ptpConformance 1}
ptpDefaultDsGrp OBJECT-GROUP
    OBJECTS {
        ptpDefaultDsTwoStepFlag,
```

```

    ptpDefaultDsClockIdentity,
    ptpDefaultDsNumberPorts,
    ptpDefaultDsClockClass,
    ptpDefaultDsClockAccuracy,
    ptpDefaultDsOffsetScaledLogVaria,
    ptpDefaultDsPriority1,
    ptpDefaultDsPriority2,
    ptpDefaultDsDomainNumber,
    ptpDefaultDsSlaveOnly,
    ptpDefaultDsProfileSet,
    ptpDefaultDsTimeInaccuracy,
    ptpDefaultDsOffsetFromMasterLim
  }
  STATUS current
  DESCRIPTION
    "Objects in the default dataset"
  ::= {ptpGroups 1}

ptpCurrentDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpCurrentDsStepsRemoved,
    ptpCurrentDsOffsetFromMaster,
    ptpCurrentDsMeanLinkDelay
  }
  STATUS current
  DESCRIPTION
    "Objects in the current dataset"
  ::= {ptpGroups 2}

ptpParentDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpParentDsClockIdentity,
    ptpParentDsPortNumber,
    ptpParentDsStats,
    ptpParentDsObsOfstScdLVar,
    ptpParentDsObsPhChgRate,
    ptpParentDsGmIdentity,
    ptpParentDsGmClass,
    ptpParentDsGmAccuracy,
    ptpParentDsGmOfstScdLVar,
    ptpParentDsGmPriority1,
    ptpParentDsGmPriority2
  }
  STATUS current
  DESCRIPTION
    "Objects in the parent dataset"
  ::= {ptpGroups 3}

ptpTimePropDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpTimePropDsCurrentUtcOffset,
    ptpTimePropDsCurrentUtcOffsetVld,
    ptpTimePropDsLeap59,
    ptpTimePropDsLeap61,
    ptpTimePropDsTimeTraceable,
    ptpTimePropDsFrequencyTraceable,
    ptpTimePropDsPtpTimescale,
    ptpTimePropDsTimeSource
  }
  STATUS current
  DESCRIPTION
    "Objects in the time properties dataset"
  ::= {ptpGroups 4}

ptpForeignMastersDsGrp  OBJECT-GROUP
  OBJECTS {

```

```

    ptpForeignMasterDsPortIdentity,
    ptpForeignMasterAnnounceMessages
}
STATUS current
DESCRIPTION
    "Objects in the foreign master dataset"
::= {ptpGroups 5}

ptpPortDsGrp OBJECT-GROUP
OBJECTS {
    ptpPortDsClockIdentity,
    ptpPortDsPortNumber,
    ptpPortDsPortState,
    ptpPortDsLogMinDelayReqInterval,
    ptpPortDsPeerMeanLinkDelay,
    ptpPortDsLogAnnounceInterval,
    ptpPortDsAnnounceReceiptTimeout,
    ptpPortDsLogSyncInterval,
    ptpPortDsDelayMechanism,
    ptpPortDsLogMinPdelayReqInterval,
    ptpPortDsVersionNumber,
    ptpPortDsNetworkProtocol,
    ptpPortDsPortEnabled,
    ptpPortDsDlyAsymmetry,
    ptpPortDsProfileId,
    ptpPortDsVlanEnable,
    ptpPortDsVlanId,
    ptpPortDsVlanPrio,
    ptpPortDsTwoStepFlag,
    ptpPortDsPeerIdentity,
    ptpPortDsPrpAttachement,
    ptpPortDsPrpPairedPort,
    ptpPortDsErrorCounter,
    ptpPortDsPeerDelayLim
}
STATUS current
DESCRIPTION
    "Objects in the port dataset"
::= {ptpGroups 6}

ptpTcDefaultDsGrp OBJECT-GROUP
OBJECTS {
    ptpTcDefaultDsClkIdentity,
    ptpTcDefaultDsNumberPorts,
    ptpTcDefaultDsDelayMech,
    ptpTcDefaultDsPriDomain,
    ptpTcDefaultDsProfileSet,
    ptpTcDefaultDsTimeInaccuracy
}
STATUS current
DESCRIPTION
    "Objects in the transparent clock default dataset"
::= {ptpGroups 7}

ptpTcPortDsGrp OBJECT-GROUP
OBJECTS {
    ptpTcPortDsClockIdentity,
    ptpTcPortDsPortNumber,
    ptpTcPortDsLogMinPdelayReqInterv,
    ptpTcPortDsFaultyFlag,
    ptpTcPortDsPeerMeanLinkDelay,
    ptpTcPortDsPortEnabled,
    ptpTcPortDsDlyAsymmetry,
    ptpTcPortDsTwoStepFlag,
    ptpTcPortDsPeerIdentity,
    ptpTcPortDsPrpAttachement,
    ptpTcPortDsPrpPairedPort,

```

```

        ptpTcPortDsErrorCounter,
        ptpTcPortDsPeerDelayLim
    }
    STATUS current
    DESCRIPTION
        "Objects in the transparent clock port dataset"
    ::= { ptpGroups 8 }

ptpTlvGrp OBJECT-GROUP
    OBJECTS {
        ptpAtoiCurrentOffset,
        ptpAtoiJumpSeconds,
        ptpAtoiTimeOfNextJump,
        ptpAtoiDisplayName
    }
    STATUS current
    DESCRIPTION
        "Objects in the TLV dataset"
    ::= { ptpGroups 9 }

ptpNotificationGrp NOTIFICATION-GROUP
    NOTIFICATIONS {
        eventFaultyState,
        eventRedundancyLoss,
        eventPeerDelayExceeded
    }
    STATUS current
    DESCRIPTION
        "Notifications"
    ::= { ptpGroups 10 }

=====
-- MIB module Compliance statements
=====

ptpCompliances OBJECT IDENTIFIER ::= { ptpConformance 2 }
ptpCompliance MODULE-COMPLIANCE
    STATUS current
    DESCRIPTION
        "Compliance statement for support by of the IEC 61439-8 module"
    MODULE
    MANDATORY-GROUPS {
        eventPeerDelayExceeded,
        ptpDefaultDsGrp,
        ptpCurrentDsGrp,
        ptpParentDsGrp,
        ptpTimePropDsGrp,
        ptpForeignMastersDsGrp,
        ptpPortDsGrp,
        ptpTcDefaultDsGrp,
        ptpTcPortDsGrp,
        ptpTlvGrp,
        ptpNotificationGrp
    }
    ::= { ptpCompliances 1 }
END

```

Bibliography

IEC ~~61580~~ 61850 (all parts), *Communication networks and systems for power utility automation*

~~IEC 61850-90-4, Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines for substations²~~

IEC 61850-8-1, *Communication networks and systems for power utility automation – Part 8-1: Specific communication service mapping (SCSM) – Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3*

IEC 61850-9-2, *Communication networks and systems for power utility automation – Part 9-2: Specific communication service mapping (SCSM) – Sampled values over ISO/IEC 8802-3*

IEC 62439-2, *Industrial communication networks – High availability automation networks – Part 2: Media Redundancy Protocol (MRP)*

IEC 62439-6, *Industrial communication networks – High availability automation networks – Part 6: Distributed Redundancy Protocol (DRP)*

IEC 62439-7, *Industrial communication networks – High availability automation networks – Part 7: Ring-based Redundancy Protocol (RRP)*

ISO/IEC 8802-2, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 2: Logical Link Control*

ISO/IEC 9646-7, *Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 7: Implementation Conformance Statements*

RFC 4330 (2006), *The Internet Society, Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*

² ~~To be published.~~

INTERNATIONAL STANDARD

NORME INTERNATIONALE



**Industrial communication networks – High availability automation networks –
Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless
Redundancy (HSR)**

**Réseaux de communication industriels – Réseaux d'automatisme à haute
disponibilité –**

**Partie 3: Protocole de redondance en parallèle (PRP) et redondance transparente
de haute disponibilité (HSR)**

CONTENTS

FOREWORD.....	8
INTRODUCTION.....	10
0.1 General.....	10
0.2 Changes with respect to the previous edition	10
0.3 Patent declaration	10
1 Scope.....	12
2 Normative references.....	12
3 Terms, definitions, abbreviations, acronyms, and conventions	13
3.1 Terms and definitions	13
3.2 Abbreviations and acronyms.....	13
3.3 Conventions.....	14
4 Parallel Redundancy Protocol (PRP)	15
4.1 PRP principle of operation.....	15
4.1.1 PRP network topology.....	15
4.1.2 PRP LANs with linear or bus topology.....	15
4.1.3 PRP LANs with ring topology.....	16
4.1.4 DANP node structure	16
4.1.5 PRP attachment of singly attached nodes.....	17
4.1.6 Compatibility between singly and doubly attached nodes	18
4.1.7 Network management	18
4.1.8 Implication on application	18
4.1.9 Transition to non-redundant networks.....	18
4.1.10 Duplicate handling	19
4.1.11 Network supervision.....	24
4.1.12 Redundancy management interface	24
4.2 PRP protocol specifications	24
4.2.1 Installation, configuration and repair guidelines	24
4.2.2 Unicast MAC addresses	25
4.2.3 Multicast MAC addresses	25
4.2.4 IP addresses.....	25
4.2.5 Nodes.....	25
4.2.6 Duplicate Accept mode (testing only).....	26
4.2.7 Duplicate Discard mode	26
4.3 PRP_Supervision frame	30
4.3.1 PRP_Supervision frame format.....	30
4.3.2 PRP_Supervision frame contents	31
4.3.3 PRP_Supervision frame for RedBox	32
4.3.4 Reception of a PRP_Supervision frame and NodesTable	32
4.4 Bridging node.....	33
4.5 Constants	33
4.6 PRP service specification	33
5 High-availability Seamless Redundancy (HSR).....	33
5.1 HSR objectives	33
5.2 HSR principle of operation.....	34
5.2.1 Basic operation with a ring topology	34
5.2.2 DANH node structure	35

5.2.3	Topology.....	36
5.2.4	RedBox structure	44
5.3	HSR node specifications	46
5.3.1	HSR operation	46
5.3.2	DANH receiving from its link layer interface	46
5.3.3	DANH receiving from an HSR port	47
5.3.4	DANH forwarding rules.....	48
5.3.5	CoS	49
5.3.6	Clock synchronization	50
5.3.7	Deterministic medium access	50
5.4	HSR RedBox specifications	50
5.4.1	RedBox properties	50
5.4.2	RedBox receiving from interlink	50
5.4.3	RedBox forwarding on the ring	52
5.4.4	RedBox receiving from an HSR port	52
5.4.5	RedBox receiving from its link layer interface.....	54
5.4.6	Redbox ProxyNodeTable handling	54
5.4.7	RedBox CoS	54
5.4.8	RedBox clock synchronization	55
5.4.9	RedBox medium access	55
5.5	QuadBox specification.....	55
5.6	Duplicate Discard method.....	55
5.7	Frame format for HSR	55
5.7.1	Frame format for all frames	55
5.7.2	HSR_Supervision frame	56
5.8	Constants	59
5.9	HSR service specification	60
6	Protocol Implementation Conformance Statement (PICS)	61
7	PRP/HSR Management Information Base (MIB).....	62
Annex A	(normative) Clocks synchronization over redundant paths in IEC 62439-3	79
A.1	Overview.....	79
A.2	Attachment to redundant LANs by a boundary clock.....	79
A.3	Attachment to redundant LANs by doubly attached ordinary clocks.....	80
A.4	PRP mapping to PTP	82
A.4.1	Scenarios and device roles	82
A.4.2	Operation in PRP	84
A.4.3	Configuration specification	85
A.4.4	Specifications of DANP as DAC.....	86
A.4.5	Clock model of a RedBox for PTP.....	86
A.5	HSR Mapping to PTP	103
A.5.1	PTP traffic in HSR.....	103
A.5.2	HSR nodes specifications.....	106
A.5.3	Redundant clocks in HSR.....	107
A.5.4	Attachment of an MC to an external LAN	107
A.6	PRP to HSR Mapping	108
A.6.1	Connection methods	108
A.6.2	PRP-HSR connection by BC.....	108
A.6.3	PRP-HSR connection by TCs	109

A.7	Doubly attached clock model	110
A.7.1	State machine	110
A.7.2	Supervision of the port	113
A.7.3	BMCA for paired ports	114
A.7.4	Selection of the port state	115
A.8	PTP datasets for high availability	115
A.8.1	General	115
A.8.2	Data types	115
A.8.3	Datasets for ordinary or boundary clocks	116
A.8.4	Object for transparent clocks	120
Annex B (normative)	PTP profile for Power Utility Automation – Redundant clock attachment	123
B.1	Application domain	123
B.2	PTP profile specification	123
B.3	Redundant clock attachment	123
Annex C (normative)	PTP profiles for high-availability automation networks	124
C.1	Application domain	124
C.2	PTP profile specification	124
C.3	Clock types	124
C.4	Protocol specification common	125
C.5	Protocol specification for L3E2E automation profile	125
C.6	Protocol specification for L2P2P automation profile	125
C.7	Timing requirements	126
C.7.1	Measurement conditions	126
C.7.2	Network time inaccuracy	126
C.7.3	Network elements	126
C.7.4	Requirements for grandmasters	126
C.7.5	Requirements for TCs	127
C.7.6	Requirements for BCs	127
C.7.7	Requirements for media converters	127
C.7.8	Requirements for links	127
C.8	Network engineering	128
C.9	Default settings	128
C.10	Redundant clock handling	129
C.11	Protocol Implementation Conformance Statement (PICS)	130
C.11.1	Conventions	130
C.11.2	PICS	130
Annex D (informative)	Precision Time Protocol tutorial for IEC 62439-3	132
D.1	Objective	132
D.2	Precision and accuracy	132
D.3	PTP clock types	133
D.4	PTP main options	134
D.5	Layer 2 and layer 3 communication	135
D.6	1-step and 2-step correction	135
D.6.1	Time correction in TCs	135
D.6.2	2-step to 1-step translation	136
D.7	End-To-End link delay measurement	138
D.7.1	General method	138
D.7.2	End-to-End link delay measurement with 1-step clock correction	138

D.7.3	End-to-End link delay measurement with 2-step clock correction	139
D.7.4	End-to-End link delay calculation by Delay_Req/Delay_Resp	140
D.8	Peer-to-Peer link delay calculation	140
D.8.1	Peer-to-Peer link delay calculation with 1-step correction	140
D.8.2	Peer-to-Peer link delay calculation with 2-step correction	141
Annex E (normative)	Management Information base for singly and doubly attached clocks	143
Bibliography		168
Figure 1	– PRP example of general redundant network	15
Figure 2	– PRP example of redundant network as two LANs (bus topology)	16
Figure 3	– PRP example of redundant ring with SANs and DANPs	16
Figure 4	– PRP with two DANPs communicating	17
Figure 5	– PRP RedBox, transition from single to double LAN	19
Figure 6	– PRP frame extended by an RCT	20
Figure 7	– PRP VLAN-tagged frame extended by an RCT	21
Figure 8	– PRP padded frame closed by an RCT	21
Figure 9	– Duplicate Discard algorithm boundaries	22
Figure 10	– HSR example of ring configuration for multicast traffic	34
Figure 11	– HSR example of ring configuration for unicast traffic	35
Figure 12	– HSR structure of a DANH	36
Figure 13	– HSR example of topology using two independent networks	37
Figure 14	– HSR example of peer coupling of two rings	38
Figure 15	– HSR example of connected rings	39
Figure 16	– HSR example of coupling two redundant PRP LANs to a ring	40
Figure 17	– HSR example of coupling from a ring node to redundant PRP LANs	41
Figure 18	– HSR example of coupling from a ring to two PRP LANs	42
Figure 19	– HSR example of coupling three rings to one PRP LAN	43
Figure 20	– HSR example of meshed topology	44
Figure 21	– HSR structure of a RedBox	45
Figure 22	– HSR frame without a VLAN tag	55
Figure 23	– HSR frame with VLAN tag	56
Figure 24	– HSR node with management counters	60
Figure 25	– HSR RedBox with management counters	61
Figure A.1	– Doubly Attached Clock as BC (MCA is best master)	79
Figure A.2	– Doubly Attached Clock when MCA is best master	81
Figure A.3	– Doubly attached clocks when OC1 is best master	82
Figure A.4	– Elements of PRP networks	84
Figure A.5	– Connection of a master clock to an ordinary clock over PRP	85
Figure A.6	– PRP RedBox as BCs (OC3 and BC7 are best masters)	87
Figure A.7	– RedBox DABC clock model	88
Figure A.8	– PRP RedBoxes as DABC with E2E – BC7 is master	89
Figure A.9	– PRP RedBoxes as DABC with E2E – timing	90

Figure A.10 – PRP RedBoxes as DABC with P2P – OC5 is best master	91
Figure A.11 – PRP RedBoxes as DABC with P2P – timing	92
Figure A.12 – PRP RedBox as DATC with E2E –signal flow	93
Figure A.13 – PRP RedBox as DATC with E2E – timing	95
Figure A.14 – PRP RedBox as DATC with P2P	96
Figure A.15 – PRP RedBox as DATC with P2P – timing	97
Figure A.16 – PRP RedBox as SLTC with E2E	100
Figure A.17 – PRP RedBox as SLTC with E2E – timing	101
Figure A.18 – PRP RedBox as SLTC with P2P	102
Figure A.19 – HSR with one GMC	104
Figure A.20 – PTP messages sent and received by an HSR node (1-step)	105
Figure A.21 – PTP messages sent and received by an HSR node (2-step)	106
Figure A.22 – Attachment of a GMC to an HSR ring through a RedBox as TC	108
Figure A.23 – PRP to HSR coupling by BCs	109
Figure A.24 – PRP to HSR coupling by TCs	110
Figure A.25 – Port states including transitions for redundant operation	111
Figure A.26 – BMCA for redundant masters	114
Figure D.1 – Precision and accuracy example	132
Figure D.2 – Precision Time Protocol principle	133
Figure D.3 – Precision Time Protocol elements	134
Figure D.4 – Delays and time-stamping logic in TCs	135
Figure D.5 – Correction of the Sync message by 1-step and 2-step (peer-to-peer)	136
Figure D.6 – Translation from 2-step to 1-step in TCs	137
Figure D.7 – Translation from 2-step to 1-step – message view	138
Figure D.8 – End-to-end link delay measurement with 1-step clock correction	139
Figure D.9 – End-to-end delay measurement with 2-step clock correction	140
Figure D.10 – Peer-to-peer link delay measurement with 1-step clock correction	141
Figure D.11 – Peer-to-peer link delay measurement with 2-step clock correction	142
Table 1 – Duplicate discard cases	23
Table 2 – Monitoring data set	26
Table 3 – NodesTable attributes	27
Table 4 – PRP_Supervision frame with no VLAN tag	30
Table 5 – PRP_Supervision frame with (optional) VLAN tag	31
Table 6 – PRP_Supervision frame contents	32
Table 7 – PRP_Supervision TLV for Redbox	32
Table 8 – PRP constants	33
Table 9 – HSR_Supervision frame with no VLAN tag	57
Table 10 – HSR_Supervision frame with optional VLAN tag	58
Table 11 – HSR Constants	60
Table A.1 – States	112
Table A.2 – Transitions	113
Table A.3 – Variables	113

Table C.1 – PTP attributes for the Industrial Automation profile	129
Table C.2 – PICS for clocks	131

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

INTERNATIONAL ELECTROTECHNICAL COMMISSION

INDUSTRIAL COMMUNICATION NETWORKS – HIGH AVAILABILITY AUTOMATION NETWORKS –

Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

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 Publication(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. IEC collaborates closely with the International Organization for Standardization (ISO) in accordance with conditions determined by agreement between the two organizations.
- 2) The formal decisions or agreements 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.
- 3) IEC Publications have the form of recommendations for international use and are accepted by IEC National Committees in that sense. While all reasonable efforts are made to ensure that the technical content of IEC Publications is accurate, IEC 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 transparently to the maximum extent possible in their national and regional publications. Any divergence between any IEC Publication and the corresponding national or regional publication shall be clearly indicated in the latter.
- 5) IEC itself does not provide any attestation of conformity. Independent certification bodies provide conformity assessment services and, in some areas, access to IEC marks of conformity. IEC is 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 its directors, employees, servants or agents including individual experts and members of its technical committees and IEC National Committees 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 Publication or any other IEC 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.

International Standard IEC 62439-3 has been prepared by subcommittee 65C: Industrial networks, of IEC technical committee 65: Industrial-process measurement, control and automation.

This third edition cancels and replaces the second edition published in 2012. This edition constitutes a technical revision.

This edition includes the following significant technical changes with respect to the previous edition:

- a) technical corrections and extension of specifications;
- b) consideration of IEC 61588 clock synchronization with end-to-end delay measurement alongside the existing peer-to-peer delay measurement in PRP.

The text of this standard is based on the following documents:

FDIS	Report on voting
65C/834/FDIS	65C/841/RVD

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

This publication has been drafted in accordance with the ISO/IEC Directives, Part 2.

This International Standard is to be read in conjunction with IEC 62439-1.

A list of all parts in the IEC 62439 series, published under the general title *Industrial communication networks – High availability automation networks*, can be found on the IEC website.

The committee has decided that the contents of this publication will remain unchanged until the stability date indicated on the IEC web site under "<http://webstore.iec.ch>" in the data related to the specific publication. At this date, the publication will be

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

IMPORTANT – The “colour inside” logo on the cover page of this publication indicates that it contains colours which are considered to be useful for the correct understanding of its contents. Users should therefore print this publication using a colour printer.

INTRODUCTION

0.1 General

IEC 62439-3 belongs to the IEC 62439 series “*Industrial communication networks – High availability automation networks*”. It specifies the PRP and HSR seamless redundancy protocols. It was adopted by IEC TC57 WG10 as the redundancy method for demanding substation automation networks operating on layer 2 networks, according to IEC 61850-8-1 and IEC 61850-9-2.

The seamless redundancy principle has been extended to clocks operating according to the Precision Time Protocol (IEC 61588) and attached to redundant networks. Two variants are specified: L3E2E for clocks which operate on layer 3 networks with end-to-end link delay measurement (E2E) and L2P2P for clocks that operate on layer 2 with peer-to-peer link delay measurement (P2P).

0.2 Changes with respect to the previous edition

The major changes with respect to IEC 62439-3:2012 are:

- Subclause 4.1.10.3 has been rewritten to explain the calculation of the duplicate rejection for different speeds.
- Annex A has been redrafted as a general concept for doubly attached clocks applicable to end-to-end (E2E) and to peer-to-peer (P2P) link delay measurement; the principle of paired port operation has now been specified in terms of a state machine based on IEC 61588:2009.
- Annex B of IEC 62439-3:2012 has been deleted; its properties are mentioned in 5.3.7.
- Annex B (new) makes the support of redundancy mandatory for IEC/IEEE 61850-9-3 that specifies doubly attached clocks on layer 2, with peer-to-peer delay measurement.
- Annex C specifies two profiles of a precision clock for industrial automation: L3E2E for layer 3, end-to-end delay measurement and L2P2P for layer 2, peer-to-peer delay measurement.
- Annex D contains the tutorial information on IEC 61588:2009 for understanding the above annexes. It was contained in IEC 62439-3:2012 Annex A.
- Annex E (MIB) contains the SNMP Management Information Base to be used for singly and doubly attached clocks in all profiles.

0.3 Patent declaration

The International Electrotechnical Commission (IEC) draws attention to the fact that it is claimed that compliance with this document may involve the use of a patent concerning filtering of redundant frames in a network node (Siemens Aktiengesellschaft – EP 2127329, US 8184650, CN 101611615B) given in 5.2.3.3.

IEC takes no position concerning the evidence, validity and scope of this patent right.

The holder of this patent right has assured the IEC that he/she is willing to negotiate licences under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of this patent right is registered with IEC. Information may be obtained from:

Siemens Aktiengesellschaft
Oto-Hahn-Ring 6
81379 Munich, Germany

The International Electrotechnical Commission (IEC) draws attention to the fact that it is claimed that compliance with this document may involve the use of patents concerning

Reception of redundant and non-redundant frames (ABB Research Ltd – EP 1825657, US 8582426, CN 101057483, IN 254425) given in 4.2.7, concerning Identifying improper cabling of devices (ABB Technology AG – EP 2163024, US 8344736, CN 101689985) given in 4.3, concerning Critical device with increased availability (ABB Research Ltd – EP 2090950) given in 4.4, concerning Ring coupling nodes for high availability networks (ABB Research Ltd – US 8582424, EP 2327185, CN 102106121) given in 5.2.3.

IEC takes no position concerning the evidence, validity and scope of these patent rights.

The holder of these patent rights has assured the IEC that he/she is willing to negotiate licences under reasonable and non-discriminatory terms and conditions with applicants throughout the world. In this respect, the statement of the holder of these patent rights is registered with IEC. Information may be obtained from:

ABB Schweiz AG
Intellectual Property CH-IP (CH-150016-L)
Brown Boveri Strasse 6
CH-5400 Baden, Switzerland
ch-ip.patent@abb.com

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights other than those identified above. IEC shall not be held responsible for identifying any or all such patent rights.

ISO (www.iso.org/patents) and IEC (<http://patents.iec.ch>) maintain on-line data bases of patents relevant to their standards. Users are encouraged to consult the data bases for the most up to date information concerning patents.

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

INDUSTRIAL COMMUNICATION NETWORKS – HIGH AVAILABILITY AUTOMATION NETWORKS –

Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)

1 Scope

The IEC 62439 series is applicable to high-availability automation networks based on the Ethernet technology.

This part of IEC 62439 specifies two redundancy protocols designed to provide seamless recovery in case of single failure of an inter-bridge link or bridge in the network, which are based on the same scheme: parallel transmission of duplicated information.

2 Normative references

The following documents, in whole or in part, are normatively referenced in this document and are indispensable for its application. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

IEC 60050-191, *International Electrotechnical Vocabulary – Chapter 191: Dependability and quality of service*

IEC 61588:2009, *Precision clock synchronization protocol for networked measurement and control systems*

IEC TR 61850-90-4:2013, *Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines*

IEC 62439-1, *Industrial communication networks – High availability automation networks – Part 1: General concepts and calculation methods*

IEC/IEEE 61850-9-3:—, *Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation (proposed IEC 61850-9-3)*¹

ISO/IEC/IEEE 8802-3:2014, *Standard for Ethernet*

IEEE 802.1D:2004, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges*

IEEE 802.1Q:2014, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges and Virtual Bridge Local Area Network*

IETF RFC 2578, *Structure of Management Information Version 2 (SMIPv2)*

¹ To be published.

IETF RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

3 Terms, definitions, abbreviations, acronyms, and conventions

3.1 Terms and definitions

For the purposes of this document, the terms and definitions given in IEC 60050-191 and in IEC 62439-1, as well as the following, apply.

3.1.1

grandmaster-capable

ordinary clock or boundary clock that is able to take the role of a grandmaster

Note 1 to entry: A grandmaster-capable clock is not necessarily connected to a recognized time source.

3.1.2

interlink

link that connects two network hierarchies

3.1.3

network time inaccuracy

time inaccuracy evaluated or measured between a particular slave clock and the time reference signal of its grandmaster clock(s), considering the worst network topology

3.1.4

RedBox

device attaching singly attached nodes to a redundant network

3.1.5

QuadBox

quadruple port device connecting two peer HSR rings, which behaves as an HSR node in each ring and is able to filter the traffic and forward it from ring to ring

3.1.6

HSR frame

frame that carries as EtherType the HSR_ethertype

3.1.7

switching logic

hardware that transmits a frame from one port to another port, possibly providing cut-through

3.1.8

time error

deviation from the time reference used for measurement or synchronization, evaluated over a short time span

3.1.9

time inaccuracy

time error not exceeded by 99,7% of the measurements, evaluated over a series of 1 000 measurements (about 20 min) in steady state

Note 1 to entry: In this document, the words “bridge” and “bridging” are synonymous with the words “switch” and “switching” respectively, when they apply to layer 2 connectivity.

3.2 Abbreviations and acronyms

For the purposes of this document, the abbreviations and acronyms given in IEC 62439-1, as well as the following apply.

ARP	Address Resolution Protocol
BC	Boundary Clock, as defined in IEC 61588:2009
BMCA	Best Master Clock Algorithm, as defined in IEC 61588:2009
BOC	Boundary Ordinary Clock
CoS	Class of Service (layer 2 quality of service)
DAC	Doubly Attached Clock
DABC	Doubly Attached Boundary Clock
DAN	Doubly Attached Node
DANH	Doubly Attached Node implementing HSR
DANP	Doubly Attached Node implementing PRP
DATC	Doubly Attached Transparent Clock
E2E	End-to-End (delay measurement)
GMC	Grandmaster Clock
HC	Hybrid Clock
HSR	High-availability Seamless Redundancy
IP	Internet Protocol
L2P2P	Layer 2 Peer-to-Peer
L3E2E	Layer 3 End-to-End
LAN	Local Area Network
LLC	Link Layer Control
LRE	Link Redundancy Entity
MAC	Medium Access Control
MC	Master Clock
MIB	Management Information Base
OC	Ordinary Clock, as defined in IEC 61588:2009
P2P	Peer to Peer (delay measurement)
PICS	Protocol Implementation Conformance Statement
PRP	Parallel Redundancy Protocol
PT	Protocol Type
RCT	Redundancy Check Trailer
RSTP	Rapid Spanning Tree Protocol
SAN	Singly Attached Node
SAC	Singly Attached Clock
SLTC	Stateless Transparent Clock
SNMP	Simple Network Management Protocol
TC	Transparent Clock, as defined in IEC 61588:2009
TCP	Transport Control Protocol
UDP	User Data Protocol
VDAN	Virtual Doubly Attached Node (SAN as visible through a RedBox)

3.3 Conventions

This part of the IEC 62439 series follows the conventions defined in IEC 62439-1.

4 Parallel Redundancy Protocol (PRP)


4.1 PRP principle of operation

4.1.1 PRP network topology

This redundancy protocol implements redundancy in the nodes rather than in the network (see the definitions in IEC 62439-1), using Doubly Attached Nodes obeying to PRP (DANPs).

A DANP is attached to two independent Local Area Networks (LANs) of similar topology, named LAN_A and LAN_B, which operate in parallel. A source DANP sends the same frame over both LANs and a destination DANP receives it from both LANs within a certain time, consumes the first frame and discards the duplicate.

Figure 1 shows a redundant network consisting of two LANs, each of which can have any topology, e.g. tree, ring or meshed.



IEC

Figure 1 – PRP example of general redundant network


The two LANs are identical in protocol at the MAC-LLC level, but they can differ in performance and topology. Transmission delays may also be different, especially if one of the networks reconfigures itself, e.g. using RSTP, to overcome an internal failure.

The two LANs follow configuration rules that allow the network management protocols such as Address Resolution Protocol (ARP) and Simple Network Management Protocol (SNMP) to operate correctly.

The two LANs have no connection between them and are assumed to be fail-independent. Redundancy can be defeated by single points of failure, such as a common power supply or a direct connection whose failure brings both networks down. Installation guidelines in this document provide guidance to the installer to achieve fail-independence.

4.1.2 PRP LANs with linear or bus topology

As an example of a simpler topology, Figure 2 draws a PRP network as two LANs in linear topology, which may also be a bus topology.




IEC

Figure 2 – PRP example of redundant network as two LANs (bus topology)

4.1.3 PRP LANs with ring topology

The two LANs can have a ring topology, as Figure 3 shows.

Redundancy within each ring is provided by a different mechanism, for instance as defined in IEEE 802.1D or in other IEC 62439 protocols.



IEC

Figure 3 – PRP example of redundant ring with SANs and DANPs

4.1.4 DANP node structure

Each node has two ports that operate in parallel and that are attached to the same upper layers of the communication stack through the Link Redundancy Entity (LRE), as Figure 4 shows.




Figure 4 – PRP with two DANPs communicating

For the basic communication, the LRE presents toward its upper layers the same interface as a non-redundant network adapter, so the upper layers are unaware of redundancy.

The LRE has two tasks: handling of duplicates and management of redundancy.

When receiving a frame from the node's upper layers, the LRE appends to the frame a Redundancy Check Trailer (RCT) containing a sequence number and sends the frame through both its ports at nearly the same time. The two frames are nearly identical except for the LAN identifier (and the checksum).

The two frames transit through the two LANs with different delays, ideally they arrive at nearly the same time at the destination node.

When receiving frames from the network, the LRE forwards the first received frame of a pair to its node's upper layers and discards the duplicate frame (if it arrives). It removes the RCT if required.

For management of redundancy and checking of the presence of other DANPs, an LRE periodically sends PRP_Supervision frames and can evaluate the PRP_Supervision frames sent by other DANPs.

4.1.5 PRP attachment of singly attached nodes

Singly attached nodes (SANs) can be attached in two ways:

SANs can be attached directly to one LAN only. Such SANs can only communicate with other SANs on the same LAN. For instance, in Figure 1, SAN A1 can communicate with SAN A2, but not with SAN B1 or SAN B2. SANs can communicate (not redundantly) with all DANPs.

SANs can be attached over a Redundancy Box (RedBox) to both LANs, as Figure 1 shows for SAN R1 and SAN R2 (see also 4.1.9). Such SANs can communicate with all DANP and SANs, for instance SAN A1 and SAN R1 can communicate.

NOTE SANs are not aware of PRP; they are for instance off-the-shelf computers or printers.

In some applications, only availability-critical devices need a double attachment, for instance the operator workplaces, while the majority of the devices are SANs. Taking advantage of the basic infrastructure of PRP, a DANP can be attached to two different bridges of the same LAN (e.g. a ring) and use protocols different from PRP to reconfigure the network in case of failure. The DANP then behaves as a bridge according to IEEE 802.1D. For instance, the bridge element may implement the RSTP protocol or a subset of RSTP where there is no forwarding of traffic between the ports. These abilities are optional and not detailed in this document. The supported mode is specified in Clause 6.

4.1.6 Compatibility between singly and doubly attached nodes

Singly attached nodes (SAN), for instance maintenance laptops or printers that belong to one LAN, can be connected to any LAN. Bridges are always SANs. A SAN connected to one LAN cannot communicate directly to a SAN connected to the other LAN. These SANs are not aware of PRP redundancy, but DANPs generate a traffic that these SANs understand.

The condition is however that the SANs ignore the RCT in the frames, which is the case since a SAN cannot distinguish the RCT from the Ethernet padding.

Conversely, DANPs understand the traffic generated by SANs, since these do not append an RCT. They only forward one frame to their upper layers since the SAN traffic uses one LAN only.

4.1.7 Network management

A node has the same MAC address on both ports and has only one set of IP addresses assigned to that address. This makes redundancy transparent to the upper layers and especially allows ARP to work as with a SAN. Bridges in a LAN are not doubly attached devices, and therefore all managed bridges have different IP (and MAC) addresses. A network management tool is preferably a DANP that can access nodes and bridges in both LANs.

4.1.8 Implication on application

When the same frame comes from the two ports with significant time difference, a frame could, in exceptional situations, be wrongly rejected as duplicate. Therefore, there exists a mechanism to age out frames. This ageing mechanism causes in seldom cases that both duplicates are accepted. The duplicate must then be filtered out by the transport layer (e.g. TCP) or tolerated by the application. This is a normal case since a LAN does not guarantee absence of duplicates.

The ageing mechanism considers both the maximum repetition rate of the frames and the worst-case difference in propagation delay, as described in 4.1.10.3.

4.1.9 Transition to non-redundant networks

The mechanism of duplicate rejection can be implemented by a RedBox that does the transition between a SAN and the doubled LANs, as Figure 5 shows. The RedBox mimics the SANs connected behind it (called VDAN or virtual DANs) and multicasts supervision frames on their behalf. The RedBox is itself a DANP and has its own IP address for management purposes, but it may also perform application functions.




Figure 5 – PRP RedBox, transition from single to double LAN

4.1.10 Duplicate handling

4.1.10.1 Duplicate Accept mode (testing only)

The Duplicate Accept mode is used for testing purposes, to verify that duplicates are indeed discarded by the link layer and not by higher layer protocols.

In the Duplicate Accept mode, a sender is configured to send a normal frame over both ports.

In the Duplicate Accept mode, a receiver is configured to accept both frames and forward both of them (if both arrive) to its upper layers.

A node indicates its mode (Duplicate Accept/Duplicate Discard) in its supervision frames.

NOTE Upper layers also discard duplicates, but less effectively than the link layer.

4.1.10.2 Duplicate Discard mode

4.1.10.2.1 Principle

To allow the receivers to detect duplicates, the sender LRE appends a six-octet field that contains a sequence number, the Redundancy Control Trailer (RCT), to both frames it sends, as Figure 6 shows.

The receiver LRE uses the sequence number of the RCT and the source MAC address to detect duplicates. It forwards only the first frame of a pair to its upper layers. To offload the application processor, the LRE can be implemented with an independent pre-processor, an intelligent Ethernet controller or in hardware.

Keeping track of duplicates allows at the same time to improve the redundancy supervision. As an option, the receiver may remove the RCT before forwarding the (first) frame to its upper layer.

The RCT consists of the following fields:

- 16-bit sequence number (SeqNr);
- 4-bit LAN identifier (LanId);
- 12 bit frame size (LSDUsize);
- 16-bit suffix (PRPsuffix).




Figure 6 – PRP frame extended by an RCT

The LRE sends both (nearly identical) frames over both LANs.

The receiving LRE can then detect duplicates based on the RCT.

NOTE 1 This method considers that SANs also exist on the network, and that frames sent by SANs could be wrongly rejected as duplicates because they happen to have a trailing field with the correct PRP suffix, same sequence number and the same size. However, SANs send on one LAN only, therefore their source MAC address will not be the same as that of a frame on the other LAN, so that a frame from an SAN will never be discarded.

NOTE 2 Reception of a well-formed RCT is no proof that the source node is a DANP. Even if the path is wrong, a DANP accepts the frame.

4.1.10.2.2 Use of PRPsuffix

The PRPsuffix identifies PRP frames and distinguishes PRP frames from other protocols that also append a trailer to their useful data.

4.1.10.2.3 Use of LSDUsize

The 12-bit LSDUsize field contains the length of the link service data unit (LSDU) in octets. The LSDUsize field allows the receiver to distinguish PRP frames from random, non-redundant frames as an additional check.

In VLANs, frame VLAN tags can be added or removed during transit through a bridge. To make the length field independent of VLAN tagging, only the LSDU and the RCT are considered in the LSDUsize, as Figure 7 shows.




Figure 7 – PRP VLAN-tagged frame extended by an RCT

If the receiver detects that the frame is closed by the correct PRP suffix, that the 12 bits before the PRP suffix correspond to the LSDU size, and that the LAN identifier matches the identifier of the LAN of the receiving port, the frame is a candidate for discarding.

Since short frames need padding to achieve the minimum frame size of 64 octets, the sender already includes the padding to avoid scanning the frame backwards, as Figure 8 shows. The minimum frame size is extended to 70 octets since the LRE or a RedBox can remove the RCT.




Figure 8 – PRP padded frame closed by an RCT

The padding size is increased to 74 octets when a VLAN tag is used.

NOTE A VLAN-tagged frame can pass several bridges, which removes or inserts VLAN tags. If the sender observes the Ethernet rule to send a minimum frame size of 68 octets for a VLAN-tagged frame and of 64 octets for a VLAN-untagged frame, there will never be a situation in which there is padding before and after the RCT.

4.1.10.2.4 Use of LanId

The 4 bit LanId field of the RCT carries a distinct identifier for LAN_A or LAN_B, specifically the codes 1010 ("A") and 1011 ("B"). Therefore, the A-frame and the B-frame differ in one bit (and in the FCS). The receiver checks that the frame comes from the correct LAN. It does not reject a frame that comes from the wrong LAN, since this could be a legitimate frame which happens to have the PRP suffix and the length information in its last 28 bits, but it increments the error counters CntErrWrongLanA or CntErrWrongLanB since this hints at a configuration error. Since this kind of error is permanent, it is detected rapidly.

4.1.10.2.5 Use of SeqNr

The 16-bit field SeqNr is incremented by one for each frame a DANP sends. The doublet {source MAC address, sequence number} uniquely identifies copies of the same frame.

4.1.10.3 Duplicate discard algorithm and configuration

The algorithm for discarding duplicates is not specified. For instance, node tables, hash tables, fifos or sequence number tracking can be used. Whichever the algorithm, it is designed such that it never rejects a legitimate frame, while occasional acceptance of a duplicate can be tolerated. Since the 16-bit SeqNr wraps around after 65 536 frames and the

sequence numbers are not necessarily contiguous, entries older than a specified time EntryForgetTime are purged.

Figure 9 shows the conditions for discarding or accepting frames.

The time t_{skew} is the time difference between the arrivals of two copies of the same frame (from LAN A to LAN B and from LAN B to LAN A) in a particular case.

The maximum time t_{skewMax} between two copies is a network property, estimated by the network designer based on the number of bridges and the traffic for a particular application, e.g. 12 ms.

The time t_{aliasMin} is the minimum possible interval between two different, legitimate frames with the same sequence number (after 65 536 increments) in a particular application.


NOTE However, a node that reboots with an arbitrary sequence number could cause a shortening of time t_{aliasMin} , as will be detailed below.

In a 100 Mbit/s network, t_{aliasMin} is the theoretical minimum frame repetition interval of the smallest frame (64 octets) from the same source ($65\,536 \times 6,72\,\mu\text{s}$) = 440 ms.

In a 1 Gbit/s, t_{aliasMin} is 44 ms.

This is a worst case since applications rarely send minimum size frames contiguously.

EXAMPLE: An IEC 61850-9-2 device sending sampled values frames at 4,8 kHz will exhibit a $t_{\text{aliasMin}} = 13\,\text{s}$.



IEC

Figure 9 – Duplicate Discard algorithm boundaries

Considering a device, entries in the duplicates list remain recorded during a time that varies between $t_{\text{resideMin}}$ and $t_{\text{resideMax}}$. These values depend on the implementation and are adequate for a number of applications, they can depend on the traffic if the list size is bound. It is recommended to make $t_{\text{resideMin}}$ and $t_{\text{resideMax}}$ adjustable depending on the data rate.

The following cases of Figure 9 are considered in Table 1:

Table 1 – Duplicate discard cases

Case	Condition	Description
Case 1	$t_{\text{skew}} > t_{\text{resideMin}}$	The duplicate is always discarded. This situation is desirable.
Case 2	$t_{\text{resideMin}} < t_{\text{skew1}} < t_{\text{resideMax}}$	The duplicate is sometimes not discarded if the entry was purged. This situation is tolerable if infrequent.
Case 3	$t_{\text{skew1}} > t_{\text{resideMax}}$	The duplicate is never discarded, the duplicate is treated as a new frame and creates a new entry in the duplicates list. This situation never occurs if the network was properly dimensioned with $(t_{\text{skewMax}} < t_{\text{resideMax}})$.
Case 4	$t_{\text{skew2}} > t_{\text{resideMax}}$	The duplicate is never discarded, but a legitimate frame is discarded since the entry created by LAN B conflicts with the alias on LAN A. This situation never occurs if the nodes are properly designed.
Case 5	$t_{\text{skew2}} > t_{\text{aliasMin}}$	The duplicate is discarded for the wrong reason, but this does not impair operation. This situation never occurs if the nodes are properly designed.

The conditions for correct operation are as follows:

- For the network engineer:

$$t_{\text{skewMax}} < (t_{\text{aliasMin}})/2 \text{ (no aliases)}$$

- For the device engineer:

$$t_{\text{resideMin}} > t_{\text{skewMax}} \text{ (recognizes all duplicates), and}$$

$$t_{\text{resideMax}} < (t_{\text{aliasMin}})/2 \text{ (prevents unwanted discards)}$$

In practice, $t_{\text{resideMin}}$ is often given by the number of frames kept in the buffer and therefore depends on the traffic, while $t_{\text{resideMax}}$ is enforced by a timer. This value is adjustable through the management setting $\text{IreDupListResideMaxTime}$ (see Clause 7)

A node that reboots could generate aliases at a lower interval than t_{aliasMin} since the sequence number at restart is arbitrary. To allow all possible destination nodes time to flush all entries of this node from the duplicates list, a node may not reboot faster than $\text{NodeRebootInterval}$ (see Table 8). $\text{NodeRebootInterval}$ holds for all nodes in the network and therefore it is independent from the link speed.

Therefore, all nodes flush their entries before $\text{NodeRebootInterval}$. The maximum residence time of an entry regardless of the speed (and of the value of $t_{\text{resideMax}}$) is EntryForgetTime .

The default value of EntryForgetTime in Table 8 (400 ms) allows in principle a t_{skewMax} of 800 ms, assuming a t_{aliasMin} of 1 600 ms. This fits all envisioned applications.

To comply with the above equations, $t_{\text{resideMax}}$ is usually much smaller than EntryForgetTime , especially in 1Gbit/s networks.

4.1.10.4 Frame size restriction

Appending the RCT could generate oversize frames that exceed the maxValidSize foreseen by ISO/IEC/IEEE 8802-3:2014 (1 522 octets).

Hence, it is expected that every network component in the LANs supports oversize frames up to maxValidSize foreseen by ISO/IEC/IEEE 8802-3:2014 plus 6 bytes of RCT (1 528 octets).

NOTE Many bridges are dimensioned for double-VLAN-tagged (Q-in-Q) frames that have a maximum size of 1 526 octets. Most Ethernet controllers are certified up to 1 528 octets. Most bridges forward correctly frames of up to 1 536 octets.

4.1.11 Network supervision

The health status of each LAN and its attached devices (nodes and bridges) is monitored, otherwise redundancy helps little.

The receiver checks that frames are correctly received over both channels. It maintains error counters that network management can read.

To this effect, senders and receivers can maintain tables that record the last time a frame was received from another node, the time a multicast or broadcast frame was sent and other protocol information.

Supervision relies on each DANP sending periodically a PRP_Supervision frame that allows checking the integrity of the network and the presence of the nodes. At the same time, these frames allow checking which devices are DANP, the MAC addresses they use and which operating mode they support: Duplicate Accept or Duplicate Discard.

4.1.12 Redundancy management interface

Redundant devices and links are useless without network management supervising redundancy and calling for maintenance actions.

The LRE presents a network management interface that allows to track the health of each LAN, and especially to detect failures early when the error rate increases. To this effect, the LRE keeps for each port a counter of received messages and of messages received with an error.

The LAN statuses appear as SNMPv1 or SNMPv2/v3 objects, allowing the same tools to manage both nodes and bridges.

4.2 PRP protocol specifications

4.2.1 Installation, configuration and repair guidelines

4.2.1.1 General

The following guidelines apply at installation time. They do not apply to conformance testing of the devices.

4.2.1.2 LANs layout

The network shall consist of two LANs that have similar properties, i.e. each one is able to carry the traffic that would exist in the absence of redundancy.

DANP shall be attached to both LANs.

SANs that need to communicate with one another shall be attached to the same LAN or to both LANs through a RedBox.

4.2.1.3 Labelling cables

The two LANs shall use cables distinctly identified as A and B.

If colours are used to identify the cables, A should be red and B should be blue.

NOTE The intention is that persons with colour disability distinguish them.

4.2.1.4 Labelling bridges

Bridges in the two LANs should have a distinct label or colour for each A or B.

4.2.1.5 Independent operation

The layout of both LANs shall fulfil the assumption of fail-independence.

4.2.1.6 Configuration

All DANPs shall be configured with the same LifeCheckInterval (see 4.5).

4.2.2 Unicast MAC addresses

Both ports A and B of a DANP shall be configured with the same MAC address.

This MAC address shall be unique in the network.

SANs connected to one LAN only shall not have the same MAC address as another node within the whole network (LAN_A plus LAN_B).

4.2.3 Multicast MAC addresses

All nodes in the network shall be configured to operate with the same multicast address for the purpose of network supervision (see 4.3).

4.2.4 IP addresses

The IP address(es) of any node or bridge shall be unique within the whole network (LAN_A plus LAN_B).

NOTE 1 Uniqueness applies to all IP addresses of a device.

A DANP shall have the same IP address(es) when seen from either LAN_A or LAN_B.

Bridges on LAN_A and LAN_B are considered as SANs and shall have different IP addresses for the purpose of network management, even if they occupy the same location.

NOTE 2 A possible scheme is to distinguish IP addresses by one bit, e.g. allocate even IP addresses to LAN_A devices and odd IP addresses to LAN_B devices.

4.2.5 Nodes

4.2.5.1 Node types

Doubly attached nodes according to the parallel redundancy protocol (DANP) shall have two ports (portA and portB) that have the same abilities, and in particular could be used alternatively if only one LAN would be connected, portA being connected to LAN_A and portB to LAN_B.

Singly Attached Nodes (SAN) have only one port for the purpose of this protocol, no special requirements apply to them.

4.2.5.2 Labelling connectors

Subclause 4.2.5.2 applies to a DANP using two LANs of similar nature.

The connectors for each LAN shall be labelled distinctly as A and B.

When connectors are ordered vertically, LAN_A should be the upper connector and LAN_B the lower connector in the normal position of the device.

When connectors are ordered horizontally, the left connector should be the LAN_A and the right connector the LAN_B, as seen from the side where the cables or fibres are plugged.

The redundant connectors shall be independently removable and insertable.

4.2.6 Duplicate Accept mode (testing only)

4.2.6.1 Sending

A sender in the Duplicate Accept mode shall send the frame it receives from its upper layers unchanged over both its ports so that the two frames appear on the respective LANs.

4.2.6.2 Receiving

A receiver in the Duplicate Accept mode shall forward frames received from either port to its upper layers.

NOTE Duplicate Accept is a test setting not used for operation.

4.2.7 Duplicate Discard mode

4.2.7.1 Monitoring data set

A node shall maintain a monitoring data set consisting of the elements of Table 2.

Table 2 – Monitoring data set

Object name	Character	Description
SendSeq	mandatory	A 16-bit sequence counter used by this node for sending
CntErrorA and CntErrorB	optional	For each port A and B, a 32-bit counter indicating the number of erroneous frames received
CntReceivedA and CntReceivedB	optional	For each port A and B, a 32-bit counter indicating the number of frames received
CntErrWrongLanA and CntErrWrongLanB	optional	For each port A and B, a 32-bit counter indicating the number of mismatches
ActiveA and ActiveB	optional	For each port A and B, the status of the communication link
NodesTable	optional	A table of all other nodes that this node detects

4.2.7.2 NodesTable

A node may maintain a NodesTable for the purpose of network supervision. If a NodesTable is used, it shall consist of the elements given in Table 3 for each entry.

Table 3 – NodesTable attributes

Argument	Definition	Data Type
MacAddress	MAC address of the remote source node (6 octets)	OctetString6
CntReceivedA	Number of frames received from that node over LAN_A	Unsigned32
CntReceivedB	Number of frames received from that node over LAN_B	Unsigned32
CntErrWrongLanA	Number of frames that were received from that node with the wrong LAN identifier on LAN_A	Unsigned32
CntErrWrongLanB	Number of frames that were received from that node with the wrong LAN identifier on LAN_B	Unsigned32
TimeLastSeenA	Time at which the latest frame was received from that node over LAN_A	TimeTicks
TimeLastSeenB	Time at which the latest frame was received from that node over LAN_B	TimeTicks
SanA	True if the remote node is most probably a SAN accessible over portA	Boolean1
SanB	True if the remote node is most probably a SAN accessible over portB	Boolean1

A node shall clear a NodesTable entry of a given other node when the time elapsed since reception of a frame from that node over both TimeLastSeenA and TimeLastSeenB exceeds NodeForgetTime (see 4.5).

NOTE 1 The NodesTable is populated by the received frames and distinguishes SANs from DANPs. The NodesTable allows a node to identify another node as SAN on one LAN only and in this case the node could send frames without an RCT over the port of that LAN only. A node without a NodesTable always sends all frames with an RCT over both ports, also when the destination is a SAN. In applications where broadcast traffic is the bulk, the impact on network load is small.

NOTE 2 The key attribute of the NodesTable is MacAddress as received in the PRP_Supervision frame sent by a DANP or in any frame received from a SAN.

NOTE 3 Most of these attributes exist not only in one instance per physical remote node, but also as separate instances for each multi/broadcast address used by that node, and some also for each multi/broadcast address used by this (local) node.

NOTE 4 TimeTicks is a local time unit, not a global time.

4.2.7.3 Redundancy Control Trailer (RCT)

The Redundancy Control Trailer (RCT) appended to DANP frames shall consist of six octets, structured in the following way (in the order of transmission):

- a 16-bit sequence number (SeqNr) transmitted with the most significant 8 bits in the first octet, which reflects the counter SendSeq (see 4.2.7.1);
- a 4-bit LAN identifier (LanId) transmitted as the most significant 4 bits of the third octet, which carries the sequence “1010” for LAN_A and the sequence “1011” for LAN_B, respectively;
- a 12 bit LSDU size (LSDUsize) whose most significant 4 bits are transmitted in the least significant 4 bits of the third octet and the least significant 8 bits in the fourth octet, that indicates the size in octets of the LSDU starting from the end of the Protocol Type (PT) field as defined in ISO/IEC/IEEE 8802-3:2014 and IEEE 802.1Q (octet offset 12-13 without LAN header or 16-17 with VLAN header) to the RCT, excluding the PT and including possible padding and the RCT itself;
- a 16-bit PRPsuffix, set to 0x88FB.

NOTE 1 Padding inserted before the RCT is included in the LSDU size.

NOTE 2 The intention is that the LSDUsize is the same in PRP and HSR.

4.2.7.4 Sending (Duplicate Discard mode)

4.2.7.4.1 Sending

When sending a frame coming from its upper layers, a node shall:

If the destination node is not registered as a SAN in the NodesTable or if the node has no NodesTable	
	If the frame carries a IEEE 802.1Q tag and has a size smaller than 64 octets (without FCS)
	pad the frame to 64 octets as specified in ISO/IEC/IEEE 8802-3:2014
	If the frame does not carry a IEEE 802.1Q tag and has a size smaller than 60 octets (without FCS)
	pad the frame to 60 octets as specified in ISO/IEC/IEEE 8802-3:2014
	increment SendSeq (see 4.2.7.1), wrapping over through 0; append the 6-octet RCT with a LanId "A" for portA and LanId "B" for portB; append the 4-octet FCS; send the frame with LanId "A" through portA and the frame with LanId "B" through portB;
Else	
	send the frame to the port(s) where the SAN is registered;

NOTE 1 Simple nodes have no NodesTable and do not register remote nodes as DANP or SAN, they always append the RCT and send over both ports.

NOTE 2 The total padding to 70 octets (no VLAN tag) considers that the LRE or a RedBox could remove the 6-octet RC. Padding to 74 octets (with VLAN tag) considers that in addition, an edge port in a bridge could remove the 4-octet priority tag. This way, the frame given to the destination respects the minimum frame size of 64 octets and the RCT is always just before the FCS.

4.2.7.4.2 Timing

The time skew between the sending of both frames of a pair should be kept at a minimum to reduce asymmetry of the transmission delays.

In particular, frames waiting for sending should be dropped from the sending queue if the corresponding port is not operational.

4.2.7.5 Receiver (Duplicate Discard mode)

4.2.7.5.1 Receiving

On reception of a frame over either port, the receiver shall:

If the frame is erroneous,	
	increment the error counter of the respective port CntErrorsA or CntErrorsB (see 4.2.7.1)
	ignore the frame;
Else if the destination MAC address belongs to the reserved addresses in IEEE 802.1Q-2004 Table 7-10,	
	treat the frame at the link layer;
Else if the frame is a duplicate (see 4.2.7.5.2 and 4.2.7.5.3),	
	discard the frame;
Else	
	forward the frame to the upper layers.

Frames with a destination MAC address according to IEEE 802.1Q:2014, Table 8-1, Table 8-2 and Table 8-3 are treated at the link layer and are not expected to carry an RCT, even if sent by a DANP.

4.2.7.5.2 Identification of frames associated with the Duplicate Discard mode

A receiver shall identify a duplicate candidate as a valid frame:

- whose last 16 bits before the checksum match the 16-bit PRPsuffix,
- whose 12 bits before the PRPsuffix match the physical size of the LSDU (excluding the ProtocolType PT, but including the RCT) and
- whose LanId matches the LAN identity of the receiving port, 1010 (A) or 1011 (B).

4.2.7.5.3 Duplicate Discard method

The Duplicate Discard method shall not discard a frame sent as single or both frames of a pair; exceptionally, both frames of a pair may be passed to the higher protocol layers.

The Duplicate Discard method shall forget an entry after EntryForgetTime (see Table 8).

4.2.7.5.4 Transparent reception

A receiver configured for TransparentReception (default) shall not remove the RCT before transferring the frame to its upper layers.

A receiver not configured for TransparentReception shall remove the RCT on frames where it has identified the presence of the RCT.

4.2.7.5.5 Receiving and NodesTable

If the node supports a NodesTable, it shall, at reception of a frame:

If this frame is not a PRP_Supervision frame or a HSR Supervision frame and its source MAC address is not yet in the NodesTable,	
	create an entry in the NodesTable for that source MAC address assuming it is a SanA or a SanB, depending which LAN the frame arrives on;
Else if this frame is a PRP_Supervision frame	
	If its source is not yet in the NodesTable
	create an entry in the NodesTable for that source assuming DANP Duplicate Accept or Duplicate Discard according to the PRP_Supervision frame contents;
	Else
	update the source status to DANP Duplicate Accept or Duplicate Discard;
record the local time at which the frame was received in the TimeLastSeenA, respectively TimeLastSeenB fields of the NodesTable for that source;	
increment by one (wrapping through 0) the counters CntReceivedA, respectively CntReceivedB of the NodesTable for that source.	
If it receives the frame from LAN_B from a node registered as SanA, or over LAN_A from a node registered as SanB,	
	set SanA = SanB = 1 for that source;
If the LanId does not match the identity of the port,	
	increment the CntErrWrongLanA, respective CntErrWrongLanB counter of the source device.

NOTE 1 Updating SanA and SanB allows moving an SAN from LAN_A to LAN_B and vice-versa. If this happens, the DANP will send on both LANs and after NodeForgetTime it will send only on the correct LAN.

4.3 PRP_Supervision frame

Each DANP shall multicast a PRP_Supervision frame over both its ports with the format specified in Table 4 (no VLAN) and Table 5 (with VLAN) respectively, every LifeCheckInterval (see 4.5).

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0					msb		U/L		I/G							
2	PRP_DestinationAddress = multicast (01-15-4E-00-01-XX)															
4									lsb							
6					msb		U/L		0							
8	SourceAddress (MAC address of the node)															
10									lsb							
12	SupEtherType = 0x88FB															
14	SupPath				SupVersion < 64											
16	SupSequenceNumber															
18	TLV1.Type = 20 or 21								TLV1.Length = 6							
20					msb		U/L		0							
22	MacAddress (MAC address of the DANP)															
24									lsb							
26	TLV2.Type = 30								TLV2.Length = 6							
28					msb		U/L		0							
30	RedBoxMacAddress															
32									lsb							
34	TLV0.Type = 0								TLV0.Length = 0							
36																
	Padding to 70 octets (no VLAN)															
60	SeqNr															
62	LanId (0x1010 or 0x1011)				LSDUsize = 52											
64	PRPsuffix = 0x88FB															
66	FCS															
68																
70																

Table 5 – PRP_Supervision frame with (optional) VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14																
16																
18																
20																
22																
24						msb	U/L	0								
26																
28									lsb							
30																
32						msb	U/L	0								
34																
36									lsb							
38																
40																
42																
44																
46																
48																
50																
52																
54																
56																
58																
60																
62																
64																
66																
68																
70																
72																
74																

NOTE 1 Octets with offset 14 to 17 are inserted only if a VLAN tag according to IEEE 802.1Q is used.

NOTE 2 The frame has a size of 70 octets (no VLAN tag) or 74 octets (with VLAN tag) to avoid padding if an entity removes the VLAN tag or the RCT.

NOTE 3 The SeqNr, LanId and LSDUsize are reused when the frame is passed to an HSR network. The SupEtherType is the same as the PRP suffix. It is also used for HSR.

NOTE 4 MAC addresses are assumed to have a length of 6 octets.

4.3.2 PRP_Supervision frame contents

The contents of a PRP_Supervision frame are defined in Table 6.

Table 6 – PRP_Supervision frame contents

Field name	Description
PRP_DestinationAddress	Multicast address 01-15-4E-00-01-XX reserved for the PRP protocol. By default XX is "00", but if conflicts arise, XX can be configured to take any value between 0x00 and 0xFF.
SourceAddress	MAC address of the node
SupEtherType	EtherType 0x88FB reserved for the PRP protocol.
SupPath	Reserved, set to 0
SupVersion	Protocol version, set to "1" for PRP. Implementation of version X of the protocol shall interpret version >X as if they were version X, ignoring any parameters and/or flags added by the more recent version, and interpret version <=X PRP_Supervision frames exactly as specified for the version concerned.
SupSequenceNumber	Sequence number incremented by 1 for each supervision frame sent.
TLV1.Type	Operation mode, a value of 20 indicates that the node supports the Duplicate Discard or a value of 21 indicates that it implements Duplicate Accept. Other values are reserved.
TLV1.Length	Length of the following MAC address in octets.
MacAddress	MAC addresses used by both port of the node.
TLV2.Type	Not present if the node is not a RedBox
TLV2.Length	Not present if the node is not a RedBox
RedBoxMacAddress	Not present if the node is not a RedBox
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

4.3.3 PRP_Supervision frame for RedBox

A RedBox, i.e. a node acting as a proxy for one or several SANs (called VDAN or virtual DAN) shall append to the TLV1 field a further TLV2 field with the content given in Table 7:

Table 7 – PRP_Supervision TLV for Redbox

Field name	Description
TLV1.Type	Operation mode with a value of 20 to indicate that the node supports the Duplicate Discard mode, since this is the only mode a RedBox supports. Other values are reserved.
TLV1.Length	Length of the following MAC address in octets.
MacAddress	MAC address of the VDAN
TLV2.Type	RedBox identifier with a value of 30 to indicate that the node is a RedBox
TLV2.Length	Length of the following MAC address in octets.
RedBoxMacAddress	MAC address of the RedBox that acts as proxy for the other device (VDAN)
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

4.3.4 Reception of a PRP_Supervision frame and NodesTable

When receiving a PRP_Supervision frame over any LAN, a node shall create an entry in the NodesTable corresponding to the MacAddress of that source as indicated in the message body, not in the source address, with the Duplicate Accept or Duplicate Discard mode as indicated in the frame.

If a node ceases to receive PRP_Supervision frames from a source for a time longer than NodeForgetTime, but receives frames from that source over one LAN only, it shall change the status of this node to SanA, respective SanB, depending on the LAN from which frames are received.

NOTE 1 This rule allows moving a SAN between LAN_A and LAN_B, and also to obtain the operating mode for a SAN if it was first registered at sending and not at receiving, since a DANP starts by sending on both LANs.

NOTE 2 This rule allows distinguishing an SAN from a DANP in Duplicate Accept mode with one line disconnected.

4.4 Bridging node

If the “bridging node” setting is enabled, the node shall act as a bridging node for its two ports, according to an unspecified protocol, i.e. HSR, RSTP or MRP, but shall not append an RCT to the frames it sends.

NOTE 1 The bridging node setting supports attachment of a DANP to two bridges of the same LAN to implement a partial redundancy topology. There is no requirement that normal frames be bridged in case of a double failure, but implementers are free to include this feature.

NOTE 2 The presence of nodes in bridging mode is potentially dangerous for a PRP network since it breaches the assumption of fail-independence of the two LANs.

4.5 Constants

The constant parameters are shown in Table 8.

Other values may be defined at the user’s responsibility.

Table 8 – PRP constants

Constant	Description	Default value
LifeCheckInterval	Typical interval between two successive PRP_Supervision frames sent by the same node.	2 000 ms
NodeForgetTime	Time after which a node entry is cleared from the NodesTable after frames from this node cease to be received.	60 000 ms
EntryForgetTime	Maximum time an entry may reside in the duplicate table, independently from speed.	400 ms
NodeRebootInterval	Minimum time during which a node that reboots remains silent, independently from speed.	500 ms

4.6 PRP service specification

Objects in the LRE, e.g. error counters, can be accessed from an application running on the same node using direct calls or from an application on another node using a management protocol such as SNMP or a specific protocol such as IEC 61850.

All LRE objects are described in Clause 7 as a Management Information Base (MIB) conformant to SNMP (RFC 2578, RFC 3418), which is common for PRP and HSR.

5 High-availability Seamless Redundancy (HSR)

5.1 HSR objectives

Clause 5 describes the application of the PRP principles of Clause 4 to implement a High-availability Seamless Redundancy (HSR), retaining the PRP property of zero recovery time, applicable to any topology, in particular rings and rings of rings.

NOTE With respect to PRP, HSR allows to roughly halve the network infrastructure. With respect to rings based on IEEE 802.1D (RSTP), IEC 62439-2 (MRP), IEC 62439-6 (DRP) or IEC 62439-7 (RRP), the network traffic is about the same.

Nodes within the ring are restricted to be HSR-capable bridging nodes, thus avoiding the use of dedicated bridges. Singly Attached Nodes (SANs) such as laptops or printers cannot be attached directly to the ring, but need attachment through a Redundancy Box (RedBox).

5.2 HSR principle of operation

5.2.1 Basic operation with a ring topology

As in PRP, a node has two ports operated in parallel; it is a DANH (Doubly Attached Node with HSR protocol).

A simple HSR network consists of doubly attached bridging nodes, each having two ring ports, interconnected by full-duplex links, as shown in the example of Figure 10 (multicast) and Figure 11 (unicast) for a ring topology.





Figure 10 – HSR example of ring configuration for multicast traffic

A source DANH sends a frame passed from its upper layers (“C” frame), prefixes it by an HSR tag to identify frame duplicates and sends the frame over each port (“A”-frame and “B”-frame).

A destination DANH receives, in the fault-free state, two identical frames from each port within a certain interval, removes the HSR tag of the first frame before passing it to its upper layers (“D”-frame) and discards any duplicate.

The nodes support the IEEE 802.1D bridge functionality and forward frames from one port to the other, according to four rules:

- 1) a node will not forward a frame that it injected into the ring or mesh;
- 2) a node will not forward a frame for which it is the unique destination (except for special applications such as redundancy supervision);
- 3) a port will not send a frame that is a duplicate of a frame that it already sent into that same direction;
- 4) a port will (optionally) refrain from sending a frame that is a duplicate of a frame that it already received from the opposite direction (except for supervision and timing frames).



IEC

Figure 11 – HSR example of ring configuration for unicast traffic

Frames circulating in the ring carry the HSR tag inserted by the source, which contains a sequence number. The doublet {source MAC address, sequence number} uniquely identifies copies of the same frame.

NOTE 1 The maximum time skew between two frames of a pair depends on the relative position of the receiving node and of the sending node. Assuming a worst case in which each node in the ring is transmitting at the same time its own frame with the largest size of 1 536 octets (maximum length supported by the EtherType defined in ISO/IEC/IEEE 8802-2:1998), each node could introduce 125 μ s of delay at 100 Mbit/s. With 50 nodes, the time skew can exceed 6 ms.

NOTE 2 The minimum delay before a frame returns to the node that injected it depends on the ring size. Assuming a worst case in which there is no other traffic and in which each node in the ring operates in cut-through with 4 μ s residence time, in a ring of 32 nodes, a frame with the maximum size of 1 536 octets (125 μ s at 100 Mbit/s) returns to the injecting node after 124 μ s, i.e. before the injecting node finished transmitting it.

5.2.2 DANH node structure

Figure 12 shows a conceptual view of the structure of a DANH implemented in hardware; practical implementations can be different. The two HSR ports A and B and the device port C are connected by the LRE, which includes a switching matrix allowing to forward frames from one port to the other. The switching matrix allows cut-through bridging. The LRE presents to the higher layers the same interface as a standard Ethernet transceiver would do.

The input circuit checks if this node is the destination of the frame and possibly does VLAN and multicast filtering to offload the processor. The Duplicate Discard is implemented in the output queues.




Figure 12 – HSR structure of a DANH

5.2.3 Topology


5.2.3.1 Attachment of singly attached nodes

Singly attached nodes (SAN), for instance maintenance laptops or printers cannot be inserted directly into the ring since they have only one port and cannot interpret the HSR tag in the frames. SANs communicate with ring devices through a RedBox that acts as a proxy for the SANs attached to it, as shown in Figure 10 and Figure 11. The RedBox is detailed in 5.2.3.8.

Connecting non-HSR nodes to ring ports, breaking the ring, is allowed to enable configuration. Non-HSR traffic within the closed ring is supported in an optional mode.

5.2.3.2 Use of HSR with separate LANs

HSR nodes can be connected in the same way as PRP nodes. To this effect, an HSR node may be set to “no forwarding” (mode N). Unlike PRP, SANs cannot be attached directly to such a duplicated network unless they are able to interpret the HSR tag.



IEC

Figure 13 – HSR example of topology using two independent networks

5.2.3.3 Peer coupling of rings

Two HSR rings may be connected by quadruple port devices with forwarding capabilities, called QuadBoxes, as Figure 14 shows. This is advantageous when the traffic flow exceeds the capabilities of a single ring. However, transmission delays from end to end are not improved.

Although one QuadBox is sufficient to conduct the traffic in the fault-free state of the network, two QuadBoxes are used to prevent a single point of failure.

A Quadbox forwards frames over each ring as any HSR node, and passes the frames unchanged to the other ring, except if the frame can be identified as a frame not to be forwarded to the other ring. To this effect, a QuadBox is expected to filter traffic based for instance on multicast filtering or on VLAN filtering. There is no learning of MAC addresses in a QuadBox, since the learning of MAC addresses on specific ports of a QuadBox device could lead to a short break in communication, if the QuadBox that has learned an address and is forwarding network traffic, fails.

With QuadBoxes realized as single physical entities, the two interconnected rings share the same redundancy domain concerning fault tolerance. If one QuadBox breaks down, both interconnected rings are in a degraded state and cannot tolerate a further fault.

Therefore, constructing QuadBoxes in the same way as a RedBox can help keep the redundancy independent. The QuadBox then consists of two devices connected by an interlink. For this reason, the RedBox specifications include the HSR connection.




Figure 14 – HSR example of peer coupling of two rings

The presence of two QuadBoxes on the same ring causes that two copies of the same frame are transferred from the first ring to the second, each generating another two copies.


This does not cause four frames to circulate on the second ring, since, when a copy from a first QuadBox reaches the second QuadBox on the same second ring, the second QuadBox will not forward it if it already sent a copy that came from its interlink. Conversely, if the second QuadBox did not yet receive a copy from its interlink, it will forward the frame, but not the copy that comes later from the interlink.

When a QuadBox receives a frame that it itself injected into the ring or a frame that the other QuadBox inserted into the ring, it forwards it to the interlink and to its other port if it did not already sent a copy. This duplicate will be discarded at the other end of the interlink. This scheme may cause some additional traffic on the interlink, but it allows to simplify the design of the logic.

NOTE The maximum time skew between two frames of a pair is about the same as if all nodes were on the same ring.

5.2.3.4 Hierarchical ring topology

An HSR network may consist of rings connected by QuadBoxes, as Figure 15 shows.



IEC

Figure 15 – HSR example of connected rings


Although a single QuadBox is sufficient to sustain the traffic, two independent QuadBoxes are needed to avoid a single point of failure.

Some SANs are connected directly to the DANH that performs the duty of a simplified RedBox.

5.2.3.5 Connection of an HSR ring to a PRP network

An HSR may be coupled to a PRP network through two RedBoxes, one for each LAN, as Figure 16 shows. In this case, the RedBoxes are configured to support PRP traffic on the interlink and HSR traffic on the ring ports.

The sequence number from the PRP RCT is reused for the HSR tag and vice versa, to allow frame identification from one network to the other and to identify pairs and duplicates on the HSR ring, introduced by a twofold injection into the ring through the two HSR RedBoxes.



IEC

Figure 16 – HSR example of coupling two redundant PRP LANs to a ring

The HSR RedBoxes for connecting the ring to a PRP network operate identically to the HSR RedBoxes used to attach SANs described in 5.2.3.1, except that they are configured as RedBox "A" or RedBox "B" to accept PRP frames on their interlink.


In Figure 16, RedBox A and RedBox B would send the same frame (A and AB and B and BA, respectively), but if a RedBox receives the frame before it could send it itself, it refrains from sending it.

In the example of Figure 16, RedBox A will not generate an "A" frame on behalf of LAN A if it previously received the same frame as "AB" from the ring, or conversely, RedBox "B" will generate an "AB" frame if it did not previously receive an "A" frame from the ring, which is the case whenever frame "A" is not a multicast frame.

Multicast frames or unicast frames without a receiver in the ring (void arrows in Figure 16) are removed by the RedBox that inserted them into the ring, if they originated from outside the

ring. To this purpose, the frames carry a LAN identifier that is at the same time the Redbox identifier.

Figure 17 shows the same coupling when the source is within the ring.



IEC


Figure 17 – HSR example of coupling from a ring node to redundant PRP LANs

To avoid reinjection of a frame into the PRP network through the other RedBox, each HSR frame carries the identifier of the PRP network from which the frame came originally, the 4-bit PathId. RedBoxes are configured and identified by the PathId of the PRP LAN to which they are attached.

5.2.3.6 Connection of several PRP networks to an HSR ring

Up to 7 PRP networks can be connected to the same HSR ring, each being identified by a 3-bit NetId.

The two RedBoxes that connect a PRP network with an HSR ring are configured with the NetId (1 to 7) and the LanId (A=0/B=1), which form the 4-bit PathId (see Figure 18).



IEC

Figure 18 – HSR example of coupling from a ring to two PRP LANs


To prevent reinjection of frames coming from one PRP network into another PRP network or from the other LAN of the same PRP network, a RedBox only forwards from the HSR ring frames that do not carry its own PathId.

When inserting into the ring a PRP frame from LAN A or from LAN B of a PRP network with a given NetId, a RedBox inserts into the PathId of the HSR tag its own NetId and the LanId, i.e. one of "2"/"3", "4"/"5", "6"/"7", "8"/"9", "A"/"B", "C"/"D" or "E"/"F", depending if it is RedBox A or RedBox B.

Conversely, when forwarding a frame from the ring to a PRP network, a RedBox inserts the LanId "A" or "B" into the RCT, depending if it is RedBox A or RedBox B.

5.2.3.7 Connection of one PRP networks to several HSR rings

A PRP network can be connected to any number of HSR rings, but these rings cannot be connected between themselves, neither by QuadBoxes nor by another PRP network since this would create loops (see Figure 19).




IEC

Figure 19 – HSR example of coupling three rings to one PRP LAN

5.2.3.8 Meshed topology

HSR allows any kind of meshing, and provides redundancy as long as the structure is free from single point of failure. For instance, Figure 20 shows a matrix arrangement of nodes. In this case, nodes have more than two ports operated in parallel that operate like the QuadBoxes. A frame received from one port is forwarded to all other ports except the one that received it, and each port forwards the frame unless it already sent a duplicate.




IEC

Figure 20 – HSR example of meshed topology

5.2.4 RedBox structure

Figure 21 shows the general structure of a RedBox:



IEC

Figure 21 – HSR structure of a RedBox

The RedBox has an LRE that performs the duties of the HSR protocol, in particular:

- forwards the frames received from one HSR port to the other HSR port, unless the frame was sent already;
- receives frames addressed to its own upper protocols;
- prefixes the frames sent by its own upper layers with the corresponding HSR tag before sending two copies over its HSR ports.

A RedBox can be operated in one of three modes: as a SAN connection, as a PRP connection or as an HSR connection (half of a QuadBox). Depending on the mode of operation, the frame handling at the interlink interface of the RedBox differs (see 5.4.2 and 5.4.4).

The switching logic in Figure 21 can be incorporated into the RedBox, so the interlink becomes an internal connection.

A simple RedBox is present in every node, since the LRE makes a transition to a single non-HSR host. Also, it is usual to have more than one host in a node, since a port for maintenance often exists.

5.3 HSR node specifications

5.3.1 HSR operation

5.3.1.1 Node properties

A DANH shall be operable in one of the following modes, which shall be changeable at runtime using management commands:

- Mode H (mandatory, default mode): HSR-tagged forwarding – in this mode, the DANH shall insert the HSR tag on behalf of its host and shall forward the ring traffic, except for frames sent by the node itself, duplicate frames and frames for which the node is the unique destination.
- Mode N (optional): No forwarding – in this mode, the node shall behave like mode H with the exception that it shall not forward ring traffic from port to port.
- Mode T (optional): Transparent forwarding – in this mode the DANH shall remove the HSR tag before forwarding the frame to the other port and shall send a frame from the host to both ports, untagged and without discarding duplicates.
- Mode M (optional): Mixed forwarding HSR-tagged and non HSR-tagged – in this mode, the DANH shall insert the HSR tag depending on local criteria when injecting frames into the ring. HSR tagged frames from ring ports are handled according to Mode H. Non-HSR tagged frames from ring ports are handled according to IEEE 802.1D forwarding rules.
- Mode U (optional): Unicast forwarding – in this mode, the node shall behave as in Mode H, except that it shall also forward traffic for which it is the unique destination.
- Mode X (optional): No sending on counter-duplicate – in this mode, the node shall behave as in Mode H, except that a port shall not send a frame that is a duplicate of a frame that it received completely and correctly from the opposite direction.

5.3.1.2 Host sequence number

An HSR node shall maintain a sequence number for each MAC address the host uses. The sequence number is initialized with 0.

NOTE The LRE is expected to detect the MAC address of the host by listening to the frames it receives from it. The host MAC address can also be configured into the LRE. Without knowledge of the host MAC address, the LRE forwards all HSR traffic to the host, treating unicast frames as multicast.

5.3.1.3 Host multicast address

A node shall have the ability to filter out at least MulticastFilterSize (see Table 11) multicast addresses or multicast address groups. If the table is not configured, the LRE shall forward all multicast frames to the host and to the other port, duplicates excluded.

5.3.1.4 Host unicast address

A node shall have the ability to communicate its MAC address(es) to the LRE. If the MAC address is not configured, the LRE shall forward all unicast frames to the host and to the other port, duplicates excluded.

5.3.1.5 Host reboot

When a node reboots, it shall refrain from sending frames over its HSR ports during a time NodeRebootTime (see Table 11) and preferably start transmission by sending a supervision frame.

5.3.2 DANH receiving from its link layer interface

For each frame to send on behalf of its link layer interface, a source node shall:

If the node is in Mode T		
	do not modify the frame;	
	duplicate the frame, enqueue it for sending into both HSR ports	
Else if the node is in Mode M		
	If the local criteria is not met	
		insert the HSR tag and increment the sequence number depending on local criteria; duplicate the frame, enqueue it for sending into both HSR ports
	Else	
		do not insert the HSR tag; process the frame according to IEEE 802.1D rules
Else		
	insert the HSR tag with the sequence number of the host; increment the sequence number, wrapping through 0; duplicate the frame, enqueue it for sending into both HSR ports	

NOTE 1 Enqueuing means that the frame will be sent as soon as no former or higher-priority frames are in the queue and the medium is ready.

NOTE 2 Sending a non-HSR frame to both ports does not produce circulating frames since such frames will not be forwarded by the adjacent HSR node. This mechanism is intended for off-ring configuration of an HSR node through a normal PC.

5.3.3 DANH receiving from an HSR port

A node receiving a frame from one of its HSR ports shall:

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

If this frame is not HSR-tagged:		
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)	
		consume the frame and do not forward it.
	Else if the node is in Mode M	
		forward according to 802.1D forwarding rules
	Else if the node is in Mode T	
		enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.
	Else	
		enqueue the unchanged frame for passing to its link layer interface. do not forward this frame
Else (HSR-tagged frame):		
	If the node is in Mode T	
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface
	Else (node is not in Mode T)	
	If this node is a destination:	
		If this is the first occurrence of the frame over the link layer interface:
		remove the HSR tag and pass the modified frame to its link layer interface.
		Else (this is not the first occurrence of the frame over the link layer interface):
		do not pass the frame to the link layer interface.
	Else (if this node is not a destination):	
		do not pass the frame to the link layer interface.
	If this node is not the only destination or the node is in Mode U:	
		If this is the first occurrence of the frame and the node is not in Mode N:
		enqueue the unmodified frame for sending over the other port(s)
		Else (this is not the first occurrence of the frame):
		discard the frame.
	Else (if this node is not in Mode U and it is the unique destination):	
		discard the frame.

NOTE 1 It is possible that more than one duplicate arrives, especially when rings are coupled.

NOTE 2 A node accepts an HSR tagged frame also if the LanId does not correspond to the PortId and if the LSDUsize does not match the frame size.

5.3.4 DANH forwarding rules

A node shall not send over a port, a frame that is a duplicate of a frame previously sent over that port in that same direction.

A node that receives a frame shall not modify its source address or HSR tag when forwarding the frame even if it did modify the contents of the frame.

A node that detects on the base of the signal quality supervision or FCS that the frame is damaged or truncated, shall not forward it. However, if the node operating in cut-through already started forwarding and then detects that the frame is damaged or truncated, it shall

append the error sequence foreseen in ISO/IEC/IEEE 8802-3:2014, 27.3.1.2.2 and then stop transmission of that frame.

If a previously connected port is disconnected from the network, a node shall purge the port's buffer so that it cannot send an obsolete frame, and resume buffering when the port is reconnected.

For frames that originate from a DANH, LanId may be set to 0 or to 1 to indicate the sending port (A = 0, B = 1)

Cleaning the transmit queues from entries coming from this node is not necessary, since a transmit queue empties itself whether the attached link is operational or not.

If a port received a complete and correct frame, it could in principle refrain from sending a duplicate of that frame in the opposite direction, since it can be assumed that all nodes in that direction already received correctly the frame since they forwarded it. This method saves bandwidth, but there are several caveats:

- This method is not applicable to supervisory frames, to preserve integrity check of the ring in both directions.
- This method is not applicable to PTP frames (see Annex A) since the delay is different in both directions.
- This method is not applicable when Mode "U" is enabled, to allow all nodes to receive all traffic.
- This method is not effective for the traffic that is not forwarded by a unique destination.
- This method only reduces the broadcast and multicast traffic when the node is not operating in cut-through mode.
- This method introduces a variable response delay depending on the position of the node in the ring and on race conditions; it cannot be predicted which node, if any, removes the duplicate.
- This method does not guarantee that the gained bandwidth remains available for other traffic.
- This method does not improve the worst case network delay, when a frame or a node is lost. Nodes presenting this behaviour would interwork with nodes not implementing it.

The ability to refrain from sending on counter-duplicate is expressed as Mode_X in the PICS.

NOTE 1 These rules remove circulating HSR frames and open the ring, in the same way as an RSTP or similar protocol. It applies to frames originally sourced by the node and to frames circulating in case a device is removed after having sent a frame, and the ring is closed again, for instance by a mechanical bridging device or when a DANH is powered down. In a ring of 50 nodes, a frame comes back to its originator after some 6 ms.

NOTE 2 These conditions enable a node to operate either in store-and-forward or in cut-through mode. Delaying the forwarding of a frame does not affect the worst-case ring delay.

NOTE 3 The Duplicate Discard method of PRP is not a preferred method for discarding duplicates in HSR, since HSR aims at preventing duplicates from circulating.

NOTE 4 The fact that the sequence numbers of the frames sent by one source are not monotonically increasing is not a reason for discarding the frame. This observation can however be used for supervision of the network.

NOTE 5 In cut-through operation, the complete reception of the HSR tag takes approximately 2,4 µs at 100 Mbit/s, after which the node decides to forward the frame or not. In store-and-forward, the complete reception of a frame takes some 123 µs at 100 Mbit/s for the maximum size frame (1 522 octets).

5.3.5 CoS

For the operation of HSR, priorities and VLANs are not required.

An HSR node is expected, as expressed in its PICS:

- to support at least 2 levels of priority according to IEEE 802.1D (IEEE 802.1p);
- to filter VLAN traffic according to IEEE 802.1Q;
- to filter multicast traffic.

5.3.6 Clock synchronization

Handling of clocks as per IEC 61588 is described in Annex A.

5.3.7 Deterministic medium access

Deterministic medium access is not specified.

NOTE Deterministic medium access can take advantage of the precise clock synchronization, by letting all nodes start transmission of frames with time-critical data at the same time in a reserved time slot and ensuring that all other frames completed transmission before the start of that time slot. The size of a frame is announced in the HSR header.

5.4 HSR RedBox specifications

5.4.1 RedBox properties

A RedBox is a device with at least three ports, two of them being ring ports for the HSR protocol, the third port being connected to an interlink.

A RedBox shall behave as a DANH for all traffic for which it is the source or the destination.

In addition to the DANH behaviour, the behaviour of the RedBox Interlink is described in 5.4.2. Since the RedBox distinguishes, on reception from the HSR network, between its internal interfaces and its interlink, the behaviour of a RedBox upon reception from an HSR network is described in 5.4.4.

A RedBox shall have a configurable NetId that is composed of three bits as defined in 5.7.1. It is used in the HSR tag for frames that were received from the interlink port of the RedBox. An HSR RedBox connection to a PRP LAN will not forward HSR frames of the configured NetId into the PRP LAN.

The NetId of the RedBox shall not be used in the HSR tag of frames generated within the RedBox. Such frame shall use a NetId of zero in the HSR tag instead, or a ring NetId that can be configured separately, with a default value of zero.

A RedBox shall be configurable for one of three modes:

- 1) HSR-SAN: the traffic on the interlink is not HSR, not PRP;
- 2) HSR-PRP: the traffic on the interlink is PRP-tagged as "A" or "B";
- 3) HSR-HSR the traffic on the interlink is HSR-tagged.

NOTE 1 A RedBox itself implements at least a management interface.

NOTE 2 A RedBox is expected to have its own IP address, especially for configuration messages. It can be accessed over the interlink or over the HSR ports.

NOTE 3 The interlink can be an internal connection if the RedBox serves as bridge at the same time.

5.4.2 RedBox receiving from interlink

When receiving a frame from its interlink port, a RedBox shall:

If the frame carries a HSR tag:			
	If the node is in Mode T		
		remove the HSR tag enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else if the RedBox operates in HSR-HSR mode		
		If the RedBox is a destination of the frame	
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else (If this is the first occurrence of the frame at the link layer interface)
			remove the HSR tag enqueue to the link layer interface of the RedBox
		If the frame is to be injected into the ring (RedBox is not only destination, Multicast/VLAN is ok)	
			If this is not the first occurrence of the frame
			discard the frame (already sent over that port)
			Else (If this is the first occurrence of the frame)
			enqueue the unmodified frame into each HSR port
	Else (If the RedBox does not operate in HSR-HSR mode)		
		discard the frame	
Else if the frame carries a PRP RCT and the RedBox is in HSR-PRP mode			
	If the node is in Mode T		
		remove the PRP RCT enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else		
		If the source MAC address is not already registered:	
			create an entry in the ProxyNodeTable;
		If the PRP tag does not correspond to the LanId of the RedBox "A" or "B"	
			increment the error counter lreCntWrongLanIdC;
		If the RedBox is a destination of the frame	
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else
			enqueue to the link layer interface of the RedBox (with the PRP RCT)
		If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok)	
			If this is not the first occurrence of the frame at each HSR port
			discard the frame (already sent over that port)
			Else (If this is the first occurrence of the frame at each HSR port)
			reuse the PRP sequence number and the RedBox's path identifier to build the HSR tag enqueue the frame into each HSR port
Else (if the frame carries neither a HSR tag nor a PRP RCT)			
	If the node is in Mode T		

		enqueue the unmodified frame to the link layer interface of the RedBox enqueue the unmodified frame into each HSR port
	Else if the node is in Mode M	
		forward according to 802.1D forwarding rules
	Else	
		If the source MAC address is not already registered in the ProxyNodeTable:
		create an entry in the ProxyNodeTable;
		Else (If the source is already registered)
		(void)
		If the RedBox is a destination of the frame:
		enqueue to the link layer interface of the RedBox
		If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok))
		append the HSR tag using the sequence number of that node; increment the sequence number of that source; enqueue the tagged frame into each HSR port

NOTE 1 Reception of an HSR frame over the interlink is considered as a configuration error in PRP connection mode. If the RedBox is used as one half of a QuadBox in HSR connection mode, then the interlink will only carry HSR traffic. In HSR-HSR mode, the interlink is handled as an additional HSR port intended for QuadBox coupling and thus does not use a ProxyNodeTable, unlike PRP-HSR and SAN-HSR RedBoxes. The RedBox implements a NodesTable to record the presence of HSR nodes.

NOTE 2 The RCT is only removed if the DANH identifies positively the sender as a DANP if the node has a NodeTable. Nodes without a NodeTable remove the RCT according to the mode setting.

5.4.3 RedBox forwarding on the ring

In addition to the forwarding rules of 5.3.4, a RedBox shall not forward in the ring a unicast frame that is intended for one of the nodes that are registered in the ProxyNodeTable. This condition is enabled by default and can be disabled by setting the RedBox to Mode U.

5.4.4 RedBox receiving from an HSR port

A RedBox that receives a valid frame over one HSR port shall:

If this frame is not HSR-tagged:			
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)		
		consume the frame and do not forward it.	
	Else if the node is in Mode M		
		enqueue the unchanged frame for passing to the interlink. forward according to 802.1D forwarding rules	
	Else if the node is in Mode T		
		enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.	
	Else		
		enqueue the unchanged frame for passing to its link layer interface do not forward this frame	
Else (frame is HSR-tagged):			
	If the node is in Mode T		
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface	
	Else		
		If this node is a destination:	
			If this is the first occurrence of the frame over the link layer interface:
			remove the HSR tag and pass the modified frame to its link layer interface.
			Else (this is not the first occurrence of the frame over the link layer interface):
			do not pass the frame to the link layer interface.
	Else (if this node is not a destination):		
		do not pass the frame to the link layer interface.	
	If this node is not the unique destination or the node is in Mode U:		
			If this is the first occurrence of the frame and the node is not in Mode N:
			enqueue the unmodified frame for sending over the other port(s).
			Else (this is not the first occurrence of the frame or the node is in Mode N):
			do not enqueue the unchanged frame to the second HSR port.
		If this is the first occurrence of the frame in direction of the interlink:	
			If the RedBox is in SAN mode:
			remove the HSR tag;
			If Source in the ProxyNodeTable
			discard the frame
			Else
			enqueue frame for passing to the interlink.
		Else if the RedBox is in PRP mode:	
			If the netId matches that of the RedBox
			discard the frame
			Else
			remove the HSR tag and append the PRP RCT with LanId ("A" or "B") of the RedBox and reusing the HSR sequence number.

					enqueue frame for passing to the interlink.
					Else (if the RedBox is in HSR mode)
					do not modify the frame.
					enqueue frame for passing to the interlink.
					Else (If this is not the first occurrence of the frame in direction of the interlink):
					discard the frame
					Else (If this node is the only (unicast) destination) and the node is not in Mode U:
					discard the frame.

NOTE For HSR-HSR connection only, a RedBox does not check if the frame was sent by one of the nodes for which it is a proxy since it cannot distinguish if the frame could have been sent by a redundant RedBox.

5.4.5 RedBox receiving from its link layer interface

A RedBox that receives a valid frame from its link layer interface shall:

If the node is in Mode T					
					do not modify the frame;
					duplicate the frame, enqueue it for sending into both HSR ports
Else If the node is in Mode M					
					insert the HSR tag and increment the sequence number depending on local criteria forward depending on local criteria
Else					
					If the RedBox is in SAN mode:
					enqueue frame for passing to the interlink.
					Else If the RedBox is in PRP mode:
					insert the PRP RCT with the sequence number of the host; enqueue frame for passing to the interlink.
					Else (if the RedBox is in HSR mode)
					insert the HSR tag with the sequence number of the host; enqueue frame for passing to the interlink.
					duplicate the frame, enqueue it for sending into both HSR ports

5.4.6 Redbox ProxyNodeTable handling

A RedBox shall hold a ProxyNodeTable containing an entry for each represented association, and supporting at least ProxyNodeTableMaxEntries entries.

A RedBox shall purge the entry of a node in the ProxyNodeTable with a configurable time-out (default of ProxyNodeTableForgetTime) for non-receiving frames of this node.

NOTE The actual size of the ProxyNodeTable is indicated in the PICS.

5.4.7 RedBox CoS

Same as 5.3.5, except that the RedBox is expected to support more VLANs and more complex and comprehensive options for engineering traffic flow in the network on higher protocol layers, e.g. multicast filtering rules, than a simple node.

5.4.8 RedBox clock synchronization

Handling of clocks as per IEC 61588 is described in Annex A.

5.4.9 RedBox medium access

Same as 5.3.7, except that the RedBox shall be able to aggregate several nodes for which it acts as a proxy.

5.5 QuadBox specification

A Quadbox shall operate conceptually as two RedBoxes in HSR-HSR mode, back-to-back.

5.6 Duplicate Discard method

For the purpose of Duplicate Discard, a frame shall be identified by:

- its source MAC address;
- its sequence number.

NOTE In addition, it is possible to use other fields of the frame such as the checksum to aid in duplicate detection.

The Duplicate Discard method is not specified.

Any Duplicate Discard method shall be able to forget an entry identified by <Source MAC Address><Sequence number> after a time EntryForgetTime.


5.7 Frame format for HSR

5.7.1 Frame format for all frames

HSR frames are identified uniquely by their HSR tag.

The HSR tag consists of:

- 16-bit Ethertype (HSR_EtherType = 0x892F)
- 4-bit path identifier (PathId)
- 12 bit frame size (LSDUsize)
- 16-bit sequence number (SeqNr)



IEC

Figure 22 – HSR frame without a VLAN tag

The 4-bit PathId field prevents reinjection of frames coming from one PRP network to another PRP network.

The 3 most significant bits of the 4-bit path field are the NetId with values 0 for regular HSR frames, 1 to 6 for frames originating from a PRP network (up to 6 networks supported, might also be used for other HSR networks in the future), 7 is reserved.

The least significant bit of the 4-bit path field is the LanId. This bit indicates for a frame coming from a PRP network if it has been received over portA or portB of a RedBox.

For frames that originate from within the HSR ring, LanId may be set to 0 or to 1 to indicate the sending port (A = 0, B = 1).

The allocation of the 4-bit PathId field is:

- 0000 to 0001 frame from an HSR node (A/B);
- 0010 to 1111 frames from one of 7 PRP networks (A/B).

If the frame carries a VLAN tag according to IEEE 802.1Q, the HSR tag shall be inserted after the VLAN tag.

The same frame with a VLAN tag is shown in Figure 23.




Figure 23 – HSR frame with VLAN tag

The minimum frame size (including the destination MAC address and the FCS) shall be 70 octets with no VLAN tag and 74 octets with a VLAN tag. The padding shall be inserted between payload and FCS.

NOTE 1 The LSDUsize has the same definition as in PRP and includes possible padding.

NOTE 2 The reason for inserting the HSR tag after the VLAN tag is to provide faster MAC address lookup in case IVL (Independent VLAN learning as defined in IEEE 802.1Q) is used.

NOTE 3 The minimum frame size ensures that the frame will meet the ISO/IEC/IEEE 8802-3:2014 minimum frame size of 64 octets after removal of the HSR tag and of the VLAN tag.

5.7.2 HSR_Supervision frame

5.7.2.1 Sending

Each DANH shall multicast an HSR_Supervision frame over both its ports with the format specified in Table 9 or Table 10.

NOTE 1 The frame has been padded to 70 octets (no VLAN) and 74 octets (with VLAN) respectively since the LRE can remove the HSR tag (6 octets) and VLAN tag (4 octets) when passing the frame to the upper layers.

Table 9 – HSR_Supervision frame with no VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0					msb		U/L		I/G							
2	HSR_DestinationAddress = multicast (01-15-4E-00-01-XX)															
4									lsb							
6					msb		U/L		0							
8	SourceAddress (MAC address of the node)															
10									lsb							
12	HSR_EtherType															
14	pathId				HSR_LSDUsize = 52 (with padding)											
16	SeqNr															
18	SupEthertype															
20	SupPath				SupVersion											
22	SupSequenceNumber															
24	TLV1.Type = 23								TLV1.Length = 6							
26					Msb		U/L		0							
28	MacAddress (MAC address of the DANH)															
30																
32	TLV2.Type = 30								TLV2.Length = 6							
34					Msb		U/L		0							
36	RedBoxMacAddress															
38									lsb							
40	TLV0.Type = 0								TLV0.Length = 0							
64	Padding to 70 octets (64 + 6 + 0 VLAN)															
66	FCS															
68																
70																

Table 10 – HSR_Supervision frame with optional VLAN tag

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14																
16																
18																
20																
22																
24																
26																
28																
30						msb	U/L	0								
32																
34																
36																
38						msb	U/L	0								
40																
42									lsb							
44																
46																
68																
70																
72																
74																

NOTE 2 The format of the frame is nearly identical to that of the PRP supervision frames.

NOTE 3 The Supervision frames allow the NodesTable to keep track of the presence of other nodes identified by their MAC address. When a node does not send a Supervision frame within a given time-out (in the order of one minute), its address is removed from the NodesTables if such are used.

5.7.2.2 HSR_Supervision frame contents

The content of the HSR_Supervision frame shall be defined as follows:

HSR_DestinationAddress	Reserved multicast address 01-15-4E-00-01-XX. By default XX is "00", but if conflicts arise, XX can be configured to take any value between 0x00 and 0xFF.
SourceAddress	MAC address of the node.
VLAN tag (optional)	IEEE 802.1Q tag if used.
HSR_EtherType	Ethertype of the HSR protocol, see 5.7.1
PathId	Path over which this frame was sent, in particular, identification of the PRP network.
LSDUsize	Size of the LSDU in octets.
SeqNr	SendSeq of the sender.
SupEtherType	Ethertype for supervision frames, shall be 0x88FB
SupPath	Unused, set to 0
SupVersion	Protocol version, set to "1" (one) for this version of HSR. Implementation of version X of the protocol shall interpret version >X as if they were version X, ignoring any parameters and/or flags added by the more recent version, and interpret version <=X exactly as specified for the version concerned. The version shall not exceed the value of 64.
SupSequenceNumber	Sequence number of the supervision frames
TLV1.Type	Has a value of 23 to indicate that it is an HSR node in normal operation. Other values are reserved.
TLV1.Length	Indicates the length of the following MAC address.
MacAddress	MAC address used by the sender.
TLV2.Type	Source node identifier set to 30 if the source node is a RedBox (see 4.3.3)
TLV2.Length	Indicates the length of the following MAC address.
RedBoxMacAddress	MAC addresses used by the RedBox on its interlink.
TLV0.Type	Closing TLV, set to 0
TLV0.Length	Closing TLV, set to 0

5.7.2.3 Reception of a HSR_Supervision frame and NodesTable

When receiving a first HSR_Supervision frame over any ring port, a node shall create an entry in the NodesTable corresponding to the MacAddress of that source as indicated in the message body, not in the source address of the frame, and register the port over which this frame came.

Subsequent reception of frames shall be registered with the objective of identifying reception errors.

5.7.2.4 Non-Reception of a HSR_Supervision frame

When a device detects another device as non-existent, it shall, if this device resumes transmission, accept frames from that device as if that device had not been connected before.

NOTE The mechanism to detect non-existing devices is not specified, it can be assumed that reading the NodesTable resets or reads a frame counter.

5.8 Constants

The constant parameters are shown in Table 11.

Other values are allowed at the user's responsibility.

Table 11 – HSR Constants

Constant	Description	Default value
LifeCheckInterval	How often a node sends an HSR_Supervision frame	2 000 ms
NodeForgetTime	Time after which a node entry is cleared in the NodesTable	60 000 ms
ProxyNodeTableForgetTime	Time after which a node entry is cleared in the ProxyNodeTable	60 000 ms
ProxyNodeTableMaxEntries	Maximum number of entries in the ProxyNodeTable	512
EntryForgetTime	Time after which an entry is removed from the duplicates	400 ms
NodeRebootInterval	Minimum time during which a node that reboots remains silent	500 ms
MulticastFilterSize	Number of multicast addresses to be filtered	16

5.9 HSR service specification

Objects in the LRE, e.g. error counters, can be accessed from an application running on the same node using direct calls or from an application on another node using a management protocol such as SNMP or a specific protocol such as IEC 61850.

All LRE objects are described in Clause 7 as an MIB conformant to SNMP (RFC 2578, RFC 3418), which is common for PRP and HSR.

For better understanding of the MIB counters, Figure 24 shows the structure of a node with the entities that maintain the counters.




Figure 24 – HSR node with management counters

Figure 25 shows the management information base for a RedBox.




Figure 25 – HSR RedBox with management counters

6 Protocol Implementation Conformance Statement (PICS)

The PICS shall indicate if the following options are supported:

- **SNMP_MIB:** ability to support the SNMP MIB
- **NTAB_SIZ:** number of entries in the NodesTable (0 = no NodesTable)
- **PRIQ_QTY:** number of supported priorities
- **VLAN_QTY:** number of supported VLANs
- **MULT_QTY:** number of filtered multicast addresses
- **CLK_1588:** support of IEC 61588 synchronization
- **PRP_SRP:** ability to perform as non-bridging node with no PRP
- **PRP_RSTP:** ability to perform as an RSTP bridge element with designated port role
- **PRP_MRP:** ability to perform as an MRP bridge element (client or master)
- **HSR_H:** ability to support HSR Mode H (default)
- **HSR_N:** ability to support HSR Mode N (no forwarding)
- **HSR_M:** ability to support HSR Mode M (mixed forwarding)
- **HSR_T:** ability to support HSR Mode T (transparent forwarding)
- **HSR_U:** ability to support HSR Mode U (forwarding also when unique destination)
- **HSR_X:** ability to support HSR Mode X (no forwarding on counter-duplicate)
- **RBX_PRP:** RedBox with PRP ports
- **RBX_HSR:** RedBox with HSR ports

- QBX_HSR: QuadBox integrating two RedBoxes
- RBX_PNT: number of entries in the ProxyNodeTable

7 PRP/HSR Management Information Base (MIB)

The MIB objects reflect the arguments of the service parameters which bear the same name, with an uppercase first letter. If the PICS option SNMP_MIB is true, the MIB data structures defined in Clause 7 shall be available at OID = 1.0.62439 in addition to the MIBs that the ports may provide, with the following definition.

NOTE When copying the MIB from a PDF document, all headers and footers are removed and characters are converted to the character format that SNMP requires. Although care has been taken not to use illegal characters, some can appear in the final document due to the editing process.

```
-- *****
IEC-62439-3-MIB DEFINITIONS ::= BEGIN
-- *****
-- Imports
-- *****
IMPORTS

    MODULE-IDENTITY, OBJECT-TYPE,
    Counter32, TimeTicks, Integer32, Unsigned32          FROM SNMPv2-SMI
    OBJECT-GROUP,
    MODULE-COMPLIANCE
    TruthValue, RowStatus, MacAddress, DisplayString,
    TEXTUAL-CONVENTION                                   FROM SNMPv2-CONF
    FROM SNMPv2-TC;

-- *****
-- Root OID
-- *****

iec62439 MODULE-IDENTITY

    LAST-UPDATED "201405220000Z" -- 2014, May 22
    ORGANIZATION "IEC/SC 65C"
    CONTACT-INFO "
        International Electrotechnical Commission
        IEC Central Office
        3, rue de Varembe
        P.O. Box 131
        CH-1211 GENEVA 20
        Switzerland
        Phone: +41 22 919 02 11
        Fax: +41 22 919 03 00
        email: info@iec.ch
    "

    DESCRIPTION "
        This MIB module defines the Network Management interfaces
        for the redundancy protocols defined by the IEC 62439 suite.
        This MIB exposes the IEC62439-3 objects (PRP + HSR)
    "

    REVISION "201405220000Z" -- 2014, May 22
    DESCRIPTION "
        added reference to ptp MIB and lreDupListResideMaxTime and type
        SecondFraction
    "

    REVISION "201202170000Z" -- February 17, 2012
    DESCRIPTION "
        Consistency brought into line to mrp, crp, brp MIBs
    "

    REVISION "201108260000Z" -- August 26, 2011
    DESCRIPTION "
```

This MIB is aligned with the changes to PRP and HSR as defined in the Amendment to IEC 62439-3

REVISION "200811100000Z" -- November 10, 2008
DESCRIPTION "

Separation of IEC 62439 into a suite of documents.
This MIB applies to IEC 62439-3, added HSR functionality

REVISION "200612160000Z" -- December 16, 2006
DESCRIPTION "

Initial version of the Network Management interface for the Parallel Redundancy Protocol

::= {iso std(0) 62439 }

-- *****
-- Redundancy Protocols
-- *****

mrp OBJECT IDENTIFIER ::= { iec62439 1 }
prp OBJECT IDENTIFIER ::= { iec62439 2 }
crp OBJECT IDENTIFIER ::= { iec62439 3 }
brp OBJECT IDENTIFIER ::= { iec62439 4 }
drp OBJECT IDENTIFIER ::= { iec62439 5 }
rrp OBJECT IDENTIFIER ::= { iec62439 6 }
ptp OBJECT IDENTIFIER ::= { iec62439 7 }

-- *****
-- Textual conventions
-- *****

SecondFraction ::= TEXTUAL-CONVENTION

DISPLAY-HINT "d"
STATUS current
DESCRIPTION

"time interval expressed in multiple of 2**⁻¹⁶ = 15,7 microseconds
this corresponds to the fraction of seconds in the NTP time representation
the minimum time is 0 microseconds, the maximum 18 hours
"

REFERENCE "IEC 62439-3"
SYNTAX Integer32

-- *****
-- Objects of the PRP Network Management
-- *****

linkRedundancyEntityNotifications OBJECT IDENTIFIER ::= { prp 20 }
linkRedundancyEntityObjects OBJECT IDENTIFIER ::= { prp 21 }
linkRedundancyEntityConformance OBJECT IDENTIFIER ::= { prp 22 }

-- *****
lreConfiguration OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 0 }
lreStatistics OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 1 }
lreConfigurationGeneralGroup OBJECT IDENTIFIER ::= { lreConfiguration 0 }
lreConfigurationInterfaceGroup OBJECT IDENTIFIER ::= { lreConfiguration 1 }
lreStatisticsInterfaceGroup OBJECT IDENTIFIER ::= { lreStatistics 1 }

-- *****
-- Objects for lreConfigurationGeneralGroup
-- *****

lreManufacturerName OBJECT-TYPE
SYNTAX DisplayString
MAX-ACCESS read-only
STATUS current
DESCRIPTION "specifies the name of the LRE device manufacturer"
::= { lreConfigurationGeneralGroup 1 }

```

lreInterfaceCount OBJECT-TYPE
    SYNTAX Integer32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION      "total number of LREs present in this system."
 ::= { lreConfigurationGeneralGroup 2 }

-- *****
-- Objects for lreConfigurationInterfacesGroup
-- *****

lreConfigurationInterfaces OBJECT IDENTIFIER
 ::= { lreConfigurationInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesConfigTable***
-- *****

lreInterfaceConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION      "
                        list of PRP/HSR LREs. Each entry corresponds
                        to one PRP/HSR Link Redundancy Entity (LRE), each representing a
                        pair of LAN ports A and B. Basic devices supporting PRP/HSR may
                        have only one LRE and thus one entry in the table, while more complex
                        devices may have several entries for multiple LREs.
                        "
 ::= { lreConfigurationInterfaces 1 }

lreInterfaceConfigEntry OBJECT-TYPE
    SYNTAX LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION      "
                        each entry contains management information applicable to a
                        particular LRE.
                        "
    INDEX { lreInterfaceConfigIndex }
 ::= { lreInterfaceConfigTable 1 }

LREInterfaceConfigEntry ::=
    SEQUENCE {
        lreInterfaceConfigIndex Unsigned32,
        lreRowStatus RowStatus,
        lreNodeType INTEGER,
        lreNodeName DisplayString,
        lreVersionName OCTET STRING,
        lreMacAddress MacAddress,
        lrePortAdminStateA INTEGER,
        lrePortAdminStateB INTEGER,
        lreLinkStatusA INTEGER,
        lreLinkStatusB INTEGER,
        lreDuplicateDiscard INTEGER,
        lreTransparentReception INTEGER,
        lreHsrLREMode INTEGER,
        lreSwitchingEndNode INTEGER,
        lreRedBoxIdentity INTEGER,
        lreEvaluateSupervision TruthValue,
        lreNodesTableClear INTEGER,
        lreProxyNodeTableClear INTEGER,
        lreDupListResideMaxTime SecondFraction
    }

lreInterfaceConfigIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible

```



```

    STATUS current
    DESCRIPTION "A unique value for each LRE."
::= { lreInterfaceConfigEntry 1 }

lreRowStatus OBJECT-TYPE
    SYNTAX RowStatus
    MAX-ACCESS read-create
    STATUS current
    DESCRIPTION "indicates the status of the LRE table entry"
::= { lreInterfaceConfigEntry 2 }

lreNodeType OBJECT-TYPE
    SYNTAX          INTEGER
                   {
prpmodel (1),
hsr (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION     "
                   specifies the operation mode of the LRE:

    PRP mode 1 (1)
    HSR mode (2)

    Note: PRP mode 0 is considered deprecated and is not supported by this
           revision of the MIB
    "
::= { lreInterfaceConfigEntry 3 }

lreNodeName OBJECT-TYPE
    SYNTAX          DisplayString
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION     "specifies this LRE's node name"
::= { lreInterfaceConfigEntry 4 }

lreVersionName OBJECT-TYPE
    SYNTAX OCTET STRING (SIZE(1..32))
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION     "specifies the version of this LRE's software"
::= { lreInterfaceConfigEntry 5 }

lreMacAddress OBJECT-TYPE
    SYNTAX MacAddress
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION     "
                   Specifies the MAC address to be used by this LRE. MAC
addresses are
                   identical for all ports of a single LRE
    "
::= { lreInterfaceConfigEntry 6 }

lrePortAdminStateA OBJECT-TYPE
    SYNTAX          INTEGER
                   {
notActive (1),
active (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION     "
                   Specifies whether the port A shall be active or not Active
through

```

```

        administrative action (Default: active).
        "
::= { lreInterfaceConfigEntry 7 }

lrePortAdminStateB OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        notActive (1),
        active (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Specifies whether the port B shall be active or not Active
through
        administrative action (Default: active).
        "
::= { lreInterfaceConfigEntry 8 }

lreLinkStatusA OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port A"
::= { lreInterfaceConfigEntry 9 }

lreLinkStatusB OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port B"
::= { lreInterfaceConfigEntry 10 }

lreDuplicateDiscard OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        doNotDiscard (1),
        discard (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        specifies whether a duplicate discard algorithm is used at
reception
        (Default: discard).
        "
::= { lreInterfaceConfigEntry 11 }

lreTransparentReception OBJECT-TYPE
    SYNTAX          INTEGER
                    {
        removeRCT (1),
        passRCT (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "

```

```

        if removeRCT is configured, the RCT is removed when
forwarding to the upper
        layers, only applicable for PRP LRE (Default: removeRCT).
    "
::= { lreInterfaceConfigEntry 12 }

lreHsrLREMode OBJECT-TYPE
    SYNTAX          INTEGER
    {
        modeh (1),
        moden (2),
        modet (3),
        modeu (4),
        modem (5)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        This enumeration is only applicable if the LRE is an HSR
bridging node or RedBox.
        It shows the mode of the HSR LRE:
        (1) Default mode: The HSR LRE is in mode h and bridges tagged HSR traffic
        (2) Optional mode: The HSR LRE is in mode n and bridging between its HSR ports
        Is disabled.
        Traffic is HSR tagged.
        (3) Optional mode: The HSR LRE is in mode t and bridges non-tagged HSR traffic
        between its HSR ports
        (4) Optional mode: The HSR LRE is in mode u and behaves like in mode h, except it
        does not remove unicast messages
        (5) Optional mode: The HSR LRE is configured in mixed mode. HSR frames are handled
        according to mode h. Non-HSR frames are handled according to
802.1D bridging rules.
        "
::= { lreInterfaceConfigEntry 13}

lreSwitchingEndNode OBJECT-TYPE
    SYNTAX          INTEGER
    {
        nonbridgingnode(1),
        bridgingunspecified(2),
        prpnode(3),
        hsrredboxsan(4),
        hsrnode(5),
        hsrredboxhsr(6),
        hsrredboxprpa(7),
        hsrredboxprpb(8)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        This enumeration shows which feature is enabled in this
particular LRE:
        (1): an unspecified non-bridging node, e.g. SRP.
        (2): an unspecified bridging node, e.g. RSTP.
        (3): a PRP node/RedBox.
        (4): an HSR RedBox with regular Ethernet traffic on its interlink.
        (5): an HSR switching node.
        (6): an HSR RedBox with HSR tagged traffic on its interlink.
        (7): an HSR RedBox with PRP traffic for LAN A on its interlink.
        (8): an HSR RedBox with PRP traffic for LAN B on its interlink.
        "
::= { lreInterfaceConfigEntry 14 }

lreRedBoxIdentity OBJECT-TYPE
    SYNTAX          INTEGER
    {
        id1a (2),
        id1b (3),

```

```

id2a (4),
id2b (5),
id3a (6),
id3b (7),
id4a (8),
id4b (9),
id5a (10),
id5b (11),
id6a (12),
id6b (13),
id7a (14),
id7b (15)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION
    "
    Applicable to RedBox HSR-PRP A and RedBox HSR-PRP B. One ID
    is used by
    one pair of RedBoxes (one configured to A and one configured to B)
    coupling an HSR ring to a PRP network. The integer value states the value
    of the path field a RedBox inserts into each frame it receives from its
    interlink and injects into the HSR ring. When interpreted as binary values,
    the LSB denotes the configuration of the RedBox (A or B), and the following
    3 bits denote the identifier of a RedBox pair.
    "
::= { lreInterfaceConfigEntry 15}

lreEvaluateSupervision OBJECT-TYPE
    SYNTAX TruthValue
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        True if the LRE evaluates received supervision frames. False
        if it
        drops the supervision frames without evaluating. Note: LREs are required
        to send supervision frames, but reception is optional. Default value is dependent
        on implementation.
        "
::= { lreInterfaceConfigEntry 16}

lreNodesTableClear OBJECT-TYPE
    SYNTAX INTEGER
    {
        noOp (0),
        clearNodeTable (1)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "specifies that the Node Table is to be cleared"
::= { lreInterfaceConfigEntry 17}

lreProxyNodeTableClear OBJECT-TYPE
    SYNTAX INTEGER
    {
        noOp (0),
        clearProxyNodeTable (1)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "specifies that the Proxy Node Table is to be cleared"
::= { lreInterfaceConfigEntry 18}

lreDupListResideMaxTime OBJECT-TYPE
    SYNTAX SecondFraction
    UNITS "binaryFractionOfSecond"
    MAX-ACCESS read-write
    STATUS current

```

DESCRIPTION "the longest time an entry may reside in the duplicates list, expressed as the number of seconds multiplied by 65536;
the default value is 26214 x 15 us, or 400 ms; too low a value can cause broadcast storms"

DEFVAL {26214}
::= { lreInterfaceConfigEntry 19}

-- *****
-- *** End lreInterfaceConfigTable ***
-- *****

-- *****
-- Objects for lreStatisticsInterfacesGroup
-- *****

lreStatisticsInterfaces OBJECT IDENTIFIER
::= { lreStatisticsInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesStatsTable ***
-- *****

lreInterfaceStatsTable OBJECT-TYPE
SYNTAX SEQUENCE OF LREInterfaceStatsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
"list of PRP/HSR LREs. Each entry corresponds to one PRP/HSR Link Redundancy Entity (LRE), each representing a pair of LAN ports A and B and a port C towards the application/interlink. Basic devices supporting PRP/HSR may have only one LRE and thus one entry in the table, while more complex devices may have several entries for multiple LREs."
::= { lreStatisticsInterfaces 1 }

lreInterfaceStatsEntry OBJECT-TYPE
SYNTAX LREInterfaceStatsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION
"An entry containing management information applicable to a particular LRE."
INDEX { lreInterfaceStatsIndex }

INDEX { lreInterfaceStatsIndex }
::= { lreInterfaceStatsTable 1 }

LREInterfaceStatsEntry ::= SEQUENCE {
lreInterfaceStatsIndex Unsigned32,
lreCntTxA Counter32,
lreCntTxB Counter32,
lreCntTxC Counter32,
lreCntErrWrongLanA Counter32,
lreCntErrWrongLanB Counter32,
lreCntErrWrongLanC Counter32,
lreCntRxA Counter32,
lreCntRxB Counter32,
lreCntRxC Counter32,
lreCntErrorsA Counter32,
lreCntErrorsB Counter32,
lreCntErrorsC Counter32,
lreCntNodes Integer32,
lreCntProxyNodes Integer32,
lreCntUniqueA Counter32,
lreCntUniqueB Counter32,

```

lreCntUniqueC Counter32,
lreCntDuplicateA Counter32,
lreCntDuplicateB Counter32,
lreCntDuplicateC Counter32,
lreCntMultiA Counter32,
lreCntMultiB Counter32,
lreCntMultiC Counter32,
lreCntOwnRxA Counter32,
lreCntOwnRxB Counter32
}

lreInterfaceStatsIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each LRE."
 ::= { lreInterfaceStatsEntry 1 }

lreCntTxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port A that are HSR tagged or
fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 2 }

lreCntTxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port B that are HSR tagged or
fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 3 }

lreCntTxC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent towards the application interface of
the DANP or DANH
        or over the interlink of the RedBox. Frames with and without PRP RCT or HSR
tag are counted, but not link-local frames.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 4 }

lreCntErrWrongLanA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "

```

number of frames with the wrong LAN identifier received on
LRE port A.

Initial value = 0. Only applicable to PRP ports.

"

::= { lreInterfaceStatsEntry 5 }

lreCntErrWrongLanB OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames with the wrong LAN identifier received on

LRE port B

Initial value = 0. Only applicable to PRP ports.

"

::= { lreInterfaceStatsEntry 6 }

lreCntErrWrongLanC OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames with the wrong LAN identifier received on

the interlink of

a RedBox. Only applicable to HSR RedBoxes in HSR-PRP configuration

(hsrredboxprpa and hsrredboxprpb).

"

::= { lreInterfaceStatsEntry 7 }

lreCntRxA OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames received on a LRE port A. Only frames that

are HSR tagged

or fitted with a PRP Redundancy Control Trailer are counted. Frames that are not forwarded anywhere (e.g. because the sender of the frame is in the proxy node table) are counted, too. Only frames received completely and without error are counted.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 8 }

lreCntRxB OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames received on a LRE port B. Only frames that

are HSR tagged

or fitted with a PRP Redundancy Control Trailer are counted. Frames that are not forwarded anywhere (e.g. because the sender of the frame is in the proxy node table) are counted, too. Only frames received completely and without error are counted.

Initial value = 0.

"

::= { lreInterfaceStatsEntry 9 }

lreCntRxC OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"

number of frames received from the application interface of a DANP or DANH or the number of number of frames received on the interlink of a RedBox. Frames with and without PRP RCT or HSR tag are counted, but not link-local frames. Only frames received completely and without error are counted.
Initial value = 0.
"

::= { lreInterfaceStatsEntry 10 }

lreCntErrorsA OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"
number of frames with errors received on this LRE port A.
Initial value = 0.
"

::= { lreInterfaceStatsEntry 11 }

lreCntErrorsB OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"
number of frames with errors received on this LRE port B.
Initial value = 0.
"

::= { lreInterfaceStatsEntry 12 }

lreCntErrorsC OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"
number of frames with errors received on the application interface of a DANP or DANH or on the interlink of a RedBox.
Initial value = 0.
"

::= { lreInterfaceStatsEntry 13 }

lreCntNodes OBJECT-TYPE

SYNTAX Integer32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"number of nodes in the Nodes Table."

::= { lreInterfaceStatsEntry 14 }

lreCntProxyNodes OBJECT-TYPE

SYNTAX Integer32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"
number of nodes in the Proxy Node Table. Only applicable to RedBox.
Initial value = 0.
"

::= { lreInterfaceStatsEntry 15 }

lreCntUniqueA OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"


```

        number of entries in the duplicate detection mechanism on
port A for which
    no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 16 }

lreCntUniqueB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 17 }

lreCntUniqueC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which no duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 18 }

lreCntDuplicateA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port A for which
    one single duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 19 }

lreCntDuplicateB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    one single duplicate was received.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 20 }

lreCntDuplicateC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which one single duplicate was
received.
        Initial value = 0.

```

```

"
::= { lreInterfaceStatsEntry 21 }

lreCntMultiA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port A for which
    more than one duplicate was received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 22 }

lreCntMultiB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on
port B for which
    more than one duplicate was received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 23 }

lreCntMultiC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of entries in the duplicate detection mechanism on the
application interface
    of the DAN or the interlink of the RedBox for which more than one duplicate was
received.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 24 }

lreCntOwnRxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of HSR tagged frames received on Port A that
originated from this
device. Frames originate from this device if the source MAC matches the
MAC of the LRE, or if the source MAC appears in the proxy node table (if
implemented). Applicable only to HSR.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 25 }

lreCntOwnRxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of HSR tagged frames received on Port B that
originated from this
device. Frames originate from this device if the source MAC matches the
MAC of the LRE, or if the source MAC appears in the proxy node table (if
implemented). Applicable only to HSR.
        Initial value = 0.
        "
::= { lreInterfaceStatsEntry 26 }

```

```
-- *****
-- *** End LRE InterfacesStatsTable ***
-- *****
```

```
-- *****
-- ***Begin LRE NodesTable ***
-- *****
```

```
lreNodesTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LRENodesEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
```

```
        "
        The node table (if it exists on that node) contains
information about        all remote LRE, which advertised themselves through
supervision frames        "
```

```
::= { lreStatisticsInterfaces 2 }
```

```
lreNodesEntry OBJECT-TYPE
    SYNTAX LRENodesEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
```

```
        "
        Each entry in the node table (if it exists) contains
information about        a particular remote LRE registered in the node table, which advertised itself
        through supervision frames.
        "
```

```
    INDEX { lreInterfaceStatsIndex, lreNodesIndex }
    ::= { lreNodesTable 1 }
```

```
LRENodesEntry::=
```

```
    SEQUENCE {
        lreNodesIndex Unsigned32,
        lreNodesMacAddress MacAddress,
        lreTimeLastSeenA TimeTicks,
        lreTimeLastSeenB TimeTicks,
        lreRemNodeType INTEGER
    }
```

```
lreNodesIndex OBJECT-TYPE
```

```
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "Unique value for each node in the LRE's node table."
```

```
::= { lreNodesEntry 1 }
```

```
lreNodesMacAddress OBJECT-TYPE
```

```
    SYNTAX MacAddress
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "Each MAC address corresponds to a single Doubly Attached Node"
```

```
::= { lreNodesEntry 2 }
```

```
lreTimeLastSeenA OBJECT-TYPE
```

```
    SYNTAX TimeTicks
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
```

```
        "
        Time in TimeTicks (1/100s) since the last frame from this
remote LRE was        received over LAN A. Initialized with a value of 0 upon node registration
```

```

        in the node table.
    "
::= { lreNodesEntry 3 }

lreTimeLastSeenB OBJECT-TYPE
    SYNTAX TimeTicks
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        Time in TimeTicks (1/100s) since the last frame from this
remote LRE was
    received over LAN B. Initialized with a value of 0 upon node registration
        in the node table.
    "
::= { lreNodesEntry 4 }

lreRemNodeType OBJECT-TYPE
    SYNTAX INTEGER
    {
        danp (0),
        redboxp (1),
        vdanp (2),
        danh (3),
        redboxh (4),
        vdanh (5)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "DAN type, as indicated in the received supervision frame"
::= { lreNodesEntry 5 }

-- *****
-- *** End LRE NodesTable ***
-- *****

-- *****
-- *** Begin LRE ProxyNodeTable ***
-- *****

lreProxyNodeTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LREProxyNodeEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        The proxy node table (if implemented) contains information
about all
        nodes, for which the LRE acts as a connection to the HSR/PRP
network.
        "
::= { lreStatisticsInterfaces 3 }

lreProxyNodeEntry OBJECT-TYPE
    SYNTAX LREProxyNodeEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION
        "
        Each entry in the proxy node table contains information about
a particular nodefor
        which the LRE acts as a connection to the HSR/PRP network.
        "
    INDEX { lreInterfaceStatsIndex, lreProxyNodeIndex }
::= { lreProxyNodeTable 1 }

LREProxyNodeEntry ::=
    SEQUENCE {
        lreProxyNodeIndex Unsigned32,
        lreProxyNodeMacAddress MacAddress
    }

```

```

    }

lreProxyNodeIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each node in the LRE's proxy node table."
    ::= { lreProxyNodeEntry 1 }

lreProxyNodeMacAddress OBJECT-TYPE
    SYNTAX MacAddress
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        Each entry contains information about a particular node
        for which the LRE acts as a proxy for the HSR/PRP network.
        "
    ::= { lreProxyNodeEntry 2 }

-- *****
-- *** End LRE ProxyNodeTable ***
-- *****

=====
-- Conformance Information
=====

linkRedundancyConformance OBJECT IDENTIFIER ::= { linkRedundancyEntityConformance 1 }
lreGroups OBJECT IDENTIFIER ::= { linkRedundancyConformance 1 }
lreDefaultGrp OBJECT-GROUP
    OBJECTS {
        lreManufacturerName,
        lreInterfaceCount,
        lreRowStatus,
        lreNodeType,
        lreNodeName,
        lreVersionName,
        lreMacAddress,
        lrePortAdminStateA,
        lrePortAdminStateB,
        lreLinkStatusA,
        lreLinkStatusB,
        lreDuplicateDiscard,
        lreTransparentReception,
        lreHsrLREMode,
        lreSwitchingEndNode,
        lreRedBoxIdentity,
        lreEvaluateSupervision,
        lreNodesTableClear,
        lreProxyNodeTableClear,
        lreDuplDistResideMaxTime,
        lreCntTxA,
        lreCntTxB,
        lreCntTxC,
        lreCntErrWrongLanA,
        lreCntErrWrongLanB,
        lreCntErrWrongLanC,
        lreCntRxA,
        lreCntRxB,
        lreCntRxC,
        lreCntErrorsA,
        lreCntErrorsB,
        lreCntErrorsC,
        lreCntNodes,
        lreCntProxyNodes,
        lreCntUniqueA,
        lreCntUniqueB,
        lreCntUniqueC,
        lreCntDuplicateA,
        lreCntDuplicateB,
    }

```

```

    lreCntDuplicateC,
    lreCntMultiA,
    lreCntMultiB,
    lreCntMultiC,
    lreCntOwnRxA,
    lreCntOwnRxB,
    lreNodesMacAddress,
    lreTimeLastSeenA,
    lreTimeLastSeenB,
    lreRemNodeType,
    lreProxyNodeMacAddress
  }
STATUS current
DESCRIPTION
  "Objects in the default group"
::= { lreGroups 1 }

-- *****
--=====
-- MIB module Compliance statements
--=====

linkRedundancyCompliances OBJECT IDENTIFIER ::= { linkRedundancyEntityConformance 2 }
linkRedundancyCompliance MODULE-COMPLIANCE
STATUS current
DESCRIPTION "Compliance for support by IEC 62439-3 module"
MODULE
MANDATORY-GROUPS {

    lreDefaultGrp
  }
::= { linkRedundancyCompliances 1 }

END

```

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

Annex A (normative)

Clocks synchronization over redundant paths in IEC 62439-3

A.1 Overview

Annex A specifies how clocks operating according to the Precision Time Protocol (PTP, IEC 61588:2009) are attached over redundant, simultaneously active paths as in PRP and HSR with no single point of failure at the network level. The use of redundant master clocks is not considered since it is already part of IEC 61588:2009 (e.g. grandmaster clusters).


This method is applied to two PTP profiles described in Annex C:

- L3E2E (Layer 3 with end-to-end delay measurement) and
- L2P2P (Layer 2 with peer-to-peer delay measurement).

Annex D contains tutorial material restricted to the aspects of IEC 61588:2009 relevant to the redundant operation in this context.

A.2 Attachment to redundant LANs by a boundary clock

An ordinary clock (OC) may be attached to redundant masters MCA and MCB by a two-port BC (see Figure A.1) which is at the same time an OC.



IEC

Figure A.1 – Doubly Attached Clock as BC (MCA is best master)

This arrangement increases the availability of both LANs in case of failure of a master. For instance in Figure A.1, if MCB fails, port B will change its state from PASSIVE_MASTER to MASTER and MCA will synchronize OC4B over the BC. If there are several BC between the two masters, only one of them will become master, according to the BMCA.

This arrangement is covered by the IEC 61588:2009 specifications. However, it requires a modification of the BMCA (see IEC 61588:2009, Figure 28) to account for the scenario in which MCA and MCB are one and the same clock. The decision for the BMCA should consider

in addition the clock quality (see Figure A.26 and A.7.3), so when port A and port B receive the same master with the same steps-removed, one of the two ports will go to PASSIVE_MASTER. This mode will not be further detailed since it is covered by IEC 61588:2009.

NOTE It has been claimed that the coupling of LAN A and LAN B by numerous BCs could introduce common failure modes.

A.3 Attachment to redundant LANs by doubly attached ordinary clocks

IEC 61588:2009 foresees that an ordinary clock has only one port. Clause A.7 extends the IEC 61588:2009 model to support an ordinary clock with redundant, paired ports, called a doubly attached clock (DAC).


The DAC principle can be applied to PRP (see Clause A.4) and to HSR (see Clause A.5). A DAC may be combined with a TC or an HC, whose closest counterpart in IEC 61588:2009 is called “*Combined ordinary and peer-to-peer transparent clock*”.

Figure A.2 shows the conceptual PTP clock model of an ordinary clock OC2 (doubly-attached slave) with two ports, one port attached to LAN A with Master Clock A (MCA) and one port attached to LAN B with Master Clock B (MCB), MCA being of better quality than both MCB and OC2. The LANs are symbolized by bridges acting as TCs.

NOTE 1 LANs can also contain BCs, in which case the BCs nearest to OC2 have the role of MCA or MCB

The redundant ports execute the PTP protocol independently. In particular, each port has a state machine according to IEC 61588:2009, 9.2.5. Each port executes the BMCA, calculates the link delays and responds to PTP management messages independently.

If a DAC has a lower clock quality than MCA or MCB, only the port connected to the better master will go to the SLAVE state and synchronize OC2, while the other port goes to the PASSIVE_SLAVE state, which does not synchronize OC2, but otherwise operates as slave port and for instance sends the delay calculation messages. The latter port will go to the SLAVE state if the other port loses synchronization or when the master on its LAN gets a better quality than the master on the other LAN.




IEC

Figure A.2 – Doubly Attached Clock when MCA is best master

NOTE 2 This concept can be extended to more than two ports or to several masters simultaneously active over the same port.

If a DAC has a better quality than both MCA and MCB, both its ports go to the MASTER state. They will send the Sync messages at approximately the same time (depending on network traffic).

Figure A.3 shows a topology in which one DAC is master (OC1) and the other is slave (OC2). Since both ports of OC2 are controlled by the same master OC1, OC2 will choose one port as SLAVE, depending on its BMCA. Since the clock identifier and quality are equal in this case, the BMCA is extended with application-specific criteria (e.g. dependent on the received clock quality) before the port identifiers are considered as a tie-breaker. A hysteresis behaviour is recommended so noisy behaviour does not cause frequent switchover.



IEC

Figure A.3 – Doubly attached clocks when OC1 is best master

NOTE 3 Such a DAC is simpler than a full BC, but this is not the main motivation. This arrangement removes the coupling between the LANs that would exist with a BC.

There are underlying dependencies between the ports. This leads to the concept of “paired ports”, which is reflected in the MIB. The dependencies are:

- The BMCA and the clock quality of OC2 is common for both ports.
- The state of one port is communicated to the other to ensure that one port does not go to the MASTER state when the other is in the SLAVE state.
- The two ports may exchange additional information, i.e. the calculated link delays.

Otherwise, when the ports are independent, the clock may send the corresponding PTP messages with the same or a different PTP sequence number on both ports.

A.4 PRP mapping to PTP

A.4.1 Scenarios and device roles


The general topology considered is shown in Figure A.4, which is not an actual topography, but only serves to explain the scenarios.

The duplicated PRP network consists of two LANs, LAN A and LAN B. The LANs are symbolized by bridges acting as TCs. BCs may exist in the LANs (not shown in Figure A.4). Ordinary clocks may be attached to the redundant LAN A and LAN B through RedBoxes as master (LAN C) and slave (LAN D).

The generalized network consists of the following elements, which are roles that a device may take, not actual devices:

- 1) OC1: DAC attached to both LANs, as master;
- 2) OC2: DAC attached to both LANs, as slave;
- 3) OC3: SAC attached to one of the PRP LANs, as master;
- 4) OC4: SAC attached to one of the PRP LANs, as slave;
- 5) OC5: SAC attached through a RedBox outside of the PRP LAN, as master;
- 6) OC6: SAC attached through a RedBox outside of the PRP LANs, as slave;
- 7) BC7: RedBox M attached to the singly attached master clock OC5, as DABC;
- 8) BC8: RedBox S attached to the singly attached slave clock OC6, as DABC;
- 9) TC9: RedBox M attached to the singly attached master clock OC5, as DATC or SLTC;
- 10) TC10: RedBox S attached to the singly attached slave clock OC6 as DATC or SLTC.

NOTE In reality, the LAN will carry a combination of all these cases, so this distinction is only conceptual.



IEC

Figure A.4 – Elements of PRP networks

A.4.2 Operation in PRP

The general PRP operating mode is that a DANP (OC1 and OC2 in Figure A.4) is attached to two distinct LANs (LAN A, LAN B in Figure A.4). For all traffic except PTP, a DANP source sends over both LANs to the same frame after appending the RCT. A frame is not modified while transiting through the LANs. The DANP destination receives the first frame of a pair and discards the duplicate on the base of its RCT and source address.


This general PRP operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same in the two LANs and some frames are modified in the TCs (bridges) while transiting through the LAN. A DANP does not receive the same PTP message from both ports even when the source is the same.

In particular:

- 1-step Syncs messages are modified by TCs to adjust the correction field.
- Follow_Up frames are generated by every 2-step clock and carry no RCT.
- TCs are not PRP-aware and not obliged to forward the RCT, which is a message part that comes after the payload.
- BCs present in the LAN are not PRP-aware and would generate their own Announce and Sync frames with no RCT appended.

Figure A.5 shows the connection between the master OC1 and a slave OC2 in a PRP network, both devices being DANPs (and DACs). The path of a Sync message through the LANs with the correction of the clock for the residence delay and link delays is suggested using a 2-step clock as an example. OC2 receives from each port the same PTP messages coming from the same master with different delays.

Since the RCT of PRP cannot be relied upon for PTP messages, ordinary clocks will not perform duplicate discard on them and will consider the two ports as independent.



IEC

Figure A.5 – Connection of a master clock to an ordinary clock over PRP

A.4.3 Configuration specification

When operating on a PRP network, an ordinary clock shall be either:

- 1) a DAC, attached directly as a DANP (see A.4.4, e.g. OC1 or OC2 in Figure A.4), in which case the clock obeys to the IEC 61588:2009 default BMCA extended by this specification;
- 2) a SAC, which is unaware of this specification, but is supposed to obey to the default BMCA IEC 61588:2009, attached either:

- to one LAN only as a singly attached clock (see OC3, OC4 in Figure A.4) without the benefit of redundancy;
- through a RedBox as specified in A.4.5 (see OC5, OC6 in Figure A.4).

NOTE A PRP topology in which one MC is located on LAN A and another MC on LAN B is considered a degraded situation, since a grandmaster clock (GMC) on one LAN ignores another master on the other LAN and cannot serve as a back-up. This situation happens when links are lost and ordinary clocks take over as master for lack of a better alternative. This situation can be avoided if at least one device according to A.7.3. exists. However, this can also be an intentional way of providing independent redundant clocks.

A.4.4 Specifications of DANP as DAC

A.4.4.1 Sending in the DACs

A DANP shall send the PTP messages over both ports according to the port state of each port, as selected by A.7.4.

NOTE There is no need to append an RCT but if one is added it would be ignored.

A.4.4.2 Reception in the DACs

A DAC shall handle all PTP messages according to A.3 for each port separately;

NOTE 1 This means in particular that a port exploits the Sync and Delay_Resp (in E2E) for supervision even if it does not use them for synchronization.

NOTE 2 A DANP does not rely on the RCT of received PTP messages.

A.4.5 Clock model of a RedBox for PTP

A.4.5.1 PRP RedBox types

RedBoxes connect singly attached devices to a duplicated LAN. There exists different kinds of RedBoxes with regard to the redundant operation. All kinds can be used and mixed in the same network. Some RedBox types introduce constraints that affect the entire network.


The following types of RedBox are considered:

- 1) RedBoxes as three-port BCs (TWBC)
The RedBox operates as a three-port BC, which acts as a slave on exactly one LAN and as a master or passive on the other two LANs. This increases resiliency of singly attached clocks on LAN A or LAN B since the BC can bridge the two LANs. This scenario is fully covered by the present IEC 61588:2009 specifications and is not repeated here (see A.4.5.2).
- 2) RedBox as a doubly attached BC (DABC)
The RedBox operates as a Doubly Attached Clock towards the duplicated LAN and as a BC towards the single LAN. The time distribution principles may be different on both sides of the BC; there is no restriction of IEC 61588:2009. (see A.4.5.3)
- 3) RedBox as a doubly-attached TC (DATC)
The RedBox operates as a DAC towards the duplicated LAN and as a TC towards the single LAN and for this it uses the BMCA to choose one path. They route the PTP messages to the correct path by manipulating the source port identity on the single-thread LAN. Restrictions to the IEC 61588:2009 are that a singly attached master receives twice as many Delay_Req as in the non-redundant case and that the number of ports in a device is restricted to 4 096 (see A.4.5.4)
- 4) RedBox as a stateless TC (SLTC)
This RedBox operates in cut-through mode and can only discard messages at reception. These RedBoxes provide higher performance for all messages. The SLTC tags PTP messages on all LANs.
The Singly Attached Clocks on the single-threaded LAN are able to operate with two simultaneous masters, as specified in IEC 61588:2009.
An SAC master on a single-threaded LAN receives twice as many Delay_Req as it would

without redundancy, causing additional processor load, and a total of four Delay_Respon would circulate on the LANs, causing additional traffic. The number of ports in a device is restricted to 4 096 (see A.4.5.5)

A.4.5.2 PRP RedBox as three-way BC (TWBC)

RedBoxes can be implemented as three-way BCs, (see A.2 and Figure A.6). This case is covered by the IEC 61588:2009 specifications.



IEC


Figure A.6 – PRP RedBox as BCs (OC3 and BC7 are best masters)

A.4.5.3 PRP RedBox as doubly attached BC (DABC)

A.4.5.3.1 RedBox as DABC – principle

A RedBox operates as a BC with two ports, one virtual port consisting of two physical ports, A and B operated as DAC towards the duplicated LAN A and LAN B and one port towards the single path LAN C. It does not operate as a BC between LAN A and LAN B.

Figure A.7 shows the conceptual PTP clock model of a RedBox as a DABC.



IEC

Figure A.7 – RedBox DABC clock model

In this configuration, two cases are distinguished:

- 1) The best master resides on the PRP LANs (OC1 in Figure A.7), port A or port B become slaves and port C competes for mastership on the non-PRP side. If it wins the BMC competition, the master OC1 will synchronize the BC, which will in turn synchronize OC6.
- 2) The best master resides on the single-threaded LAN (e.g. OC5 in Figure A.8), port C becomes slave and synchronizes the BC, whose port A and port B compete for mastership on their respective LANs, and, if successful, synchronize the DAC OC2.

A.4.5.3.2 PRP RedBox as DABC with E2E

A.4.5.3.2.1 PRP RedBox as DABC with E2E – principle

When a RedBox implements a BC with E2E (see Figure A.8), the exchange of PTP messages is limited to the individual LANs. Ordinary clocks attached to LAN A and LAN B obey to the doubly-attached clock model; they implement the E2E specified in IEC 61588:2009.

Clocks in LAN C and LAN D can use another link delay calculation (e.g. clock OC6 in Figure A.8, which uses P2P) or even a different time distribution protocol such as SNTP.





Figure A.8 – PRP RedBoxes as DABC with E2E – BC7 is master

A.4.5.3.2.2 PRP RedBox as DABC with E2E – timing diagram

Figure A.9 shows the timing of the PTP messages between the two RedBoxes. The RedBoxes behave like DACs. The PTP message exchange on the single-threaded LANs is shown in Figure A.9 with E2E as example.



IEC


Figure A.9 – PRP RedBoxes as DABC with E2E – timing

A.4.5.3.3 PRP RedBox as DABC with P2P

A.4.5.3.3.1 PRP RedBox as DABC with P2P – principle

When a RedBox implements a BC with P2P (see Figure A.10) the exchange of Sync messages is limited to the individual LANs. The P2P is link-local. The clocks attached to both LAN A and LAN B obey to the Doubly Attached Clock rules.

The time distribution protocol on LAN C or LAN D can be different, Figure A.10 shows an example of P2P on LAN C and E2E on LAN D.




IEC

Figure A.10 – PRP RedBoxes as DABC with P2P – OC5 is best master

A.4.5.3.3.2 PRP RedBox as DABC with P2P – timing diagram

Figure A.11 shows the timing diagram of the P2P, 1-step PTP message exchange between the RedBoxes. The BCs implement E2E on the single-threaded LANs to demonstrate the independence of the protocols.



IEC

Figure A.11 – PRP RedBoxes as DABC with P2P – timing

A.4.5.3.4 PRP RedBox as DABC – specification

RedBoxes acting as BCs shall behave as Doubly Attached Clocks according to A.4.4 on the PRP side.

NOTE This specification applies both to P2P and E2E. The behaviour on the single-threaded side is not specified, it is not constrained to be PTP.

A.4.5.4 PRP RedBox as Doubly Attached TC (DATC)

A.4.5.4.1 PRP RedBox as DATC with E2E

A.4.5.4.1.1 PRP RedBox as DATC with E2E – principle


All messages needed for delay measurement need to take the same path:
in 1-step, E2E, Sync and Delay_Req need to take the same path, while
in 2-step, E2E, Sync, Follow_Up, Delay_Req and Delay_Resp need to take the same path.

A DATC ensures this by letting the RedBox execute the same BMCA algorithm as a DAC would.

When the ordinary clock is a DAC (OC1, OC2), each port keeps track of the port over which it sends and receives the PTP messages, LAN A and LAN B are separate.

The most complex situation is when both the master and the slave clock are SAN attached through a RedBox, (master OC5 or slave OC6, as shown in Figure A.12). These devices

ignore over which LAN A or LAN B the PTP messages arrived. In this case, the RedBox S (TC10) makes the choice.



IEC

Figure A.12 – PRP RedBox as DATC with E2E –signal flow

The RedBox M on the master side (TC9 in Figure A.12) duplicates the Sync coming from the master clock (OC5) over LAN_C and forwards it to both port A and port B, with individual correction of the residence time. The Sync messages transit over LAN A and LAN B with different delays.

Two cases are considered (see Figure A.12):

- 1) When the slave clock is a DAC that operates with PRP (OC2), it receives both Sync and accepts one according to its BMCA. OC2 sends two Delay_Req on LAN_A and LAN_B that travel to the master OC5. The DAC OC2 will only use the Delay_Resp for time correction on the same LAN as it accepted the Sync, but the port in the PASSIVE_SLAVE state will nevertheless be able to calculate the delay on the other part, so switchover can be smoother.

- 2) When the slave clock is a SAC (OC6) that resides behind a RedBox S (TC10 in Figure A.12), TC10 receives two Syncs and selects the one from the better master according to its BMCA (A.7.4) and forwards the Sync to the slave clock (OC6) on LAN C. The slave clock OC6 responds with one Delay_Req. RedBox S (TC10) duplicates the received Delay_Req on both LANs.

The master OC5 will respond to each Delay_Req with a Delay_Resp. To ensure that the Delay_Resp takes the same path as the Delay_Req, the DATC RedBox M (TC9) manipulates the port identity of the Delay_Req, by setting the most significant two bits Bit 15 and Bit 14 to “10” for LAN A and to “11” for LAN B respectively, before forwarding it on LAN C to OC5.

The master OC5 copies the port number of the sourcePortIdentity in the Delay_Req to the port number of the requestingPortIdentity in the Delay_Resp. When the Delay_Resp returns from the master OC5, the RedBox M routes the Delay_Resp according to the manipulated port number on the same path as the Delay_Req came from. At the same time, the RedBox M sets the two most significant bits Bit 15 and Bit 14 of the port number in the requestingPortIdentity field to “00” to ensure that the slave clock recognizes its identity.

RedBox S (TC9 in Figure A.12) accepts the Delay_Resp only from the LAN over which it received the Sync (and Follow_Up in 2-step) and forwards it over LAN D; it rejects the other Delay_Resp, regardless which one came first, but registers its timely arrival to supervise the unused path in the event of a unidirectional failure.

NOTE 1 This assumes that the portNumber of a DAC on LAN A and LAN B is “00” in the two most significant bits.

NOTE 2 The RedBox also sends Delay_Req on the inactive LAN to supervise that both LANs are operating. The cost is that a master on LAN C receives twice as many Delay_Req as in the non-redundant case.

NOTE 3 This behaviour is not strictly that of a TC according to IEC 61588:2009, which does not foresee that TCs perform the BMCA. However, it is the same as for the DAC and for the BCs.

A.4.5.4.1.2 PRP RedBox as DATC with E2E –timing

The timing diagram in Figure A.13 shows the 1-step PTP message exchange between the master (OC5 in Figure A.12) and the slave (OC6 in Figure A.12) over the PRP LAN A and LAN B.

NOTE The 2-step exchange looks similar, only that a Follow_Up follows the Sync.




Figure A.13 – PRP RedBox as DATC with E2E – timing

The RedBox M on the master side (TC9 in Figure A.12) duplicates the Sync coming from the master clock (OC5) over LAN_C and forwards it to both port A and port B, with individual correction of the residence time. The Sync messages transit over LAN A and LAN B with different delays.

Assuming that the RedBox S on the slave side (TC10 in Figure A.12) considers LAN A as carrying the better master, it forwards the Sync from LAN A to LAN D, correcting its residence time. It also supervises that Sync B arrives to supervise the passive port.

When the corresponding Delay_Req arrives (same sequence number as Sync), RedBox S (TC10 in Figure A.12) forwards it only on the LAN from which it propagated the Sync to the slave.

It may send the Delay_Resp over the inactive channel to supervise that no unidirectional failure exists on the path.

The RedBox M (TC9 in Figure A.12) forwards the (both) Delay_Req to the master (OC5 in Figure A.12).


The master returns the Delay_Resp and RedBox M (TC9 in Figure A.12), forwards this response on both LANs (since it cannot know which port RedBox S prefers).

RedBox S (TC9 in Figure A.12) accepts the Delay_Resp only from the LAN over which it received the Sync (and Follow_Up in 2-step) and forwards it over LAN D; it rejects the other Delay_Resp, regardless which one came first, but registers its timely arrival to supervise the unused path in the event of a unidirectional failure.

A.4.5.4.2 PRP RedBox as DATC with P2P

A.4.5.4.2.1 PRP RedBox as DATC with P2P – principle

A RedBox can be implemented as a DATC as the model of Figure A.14 shows.



IEC

Figure A.14 – PRP RedBox as DATC with P2P

The RedBox M on the master side (TC9 in Figure A.14) duplicates the Sync.

The RedBox S on the slave side (TC10 in Figure A.14) only forwards one of the Sync.

In P2P, it does not matter if the Sync travels over LAN A and LAN B: the corrections are per TC. But in 2-step, the Follow_Up needs to take the same way as the Sync. Some clock profiles require that Announce takes the same path as the Sync.

The Sync that is not forwarded by RedBox S is nevertheless used as a health check for the path.

A.4.5.4.2.2 PRP RedBox as DATC with P2P – timing

Figure A.15 shows the timing diagram of the RedBoxes as DATCs using P2P. LAN A is assumed to contain the preferred master. Sync B is transmitted as a health check of PTP synchronization on LAN B.




Figure A.15 – PRP RedBox as DATC with P2P – timing

NOTE The timing diagram for 2-step is similar, except for the Follow_Up that follow the Sync and the Pdelay_Resp_Follow_Up that follow the Pdelay_Resp.

A.4.5.4.3 PRP RedBox as DATC – specification

A RedBox operating as a DATC shall act as a TC according to IEC 61588:2009, operate as DAC towards the PRP network and implement the paired port selection of A.7.4.

A RedBox operating as a DATC, when in E2E mode shall:

- when it receives from LAN A or LAN B a PTP message, except Delay_Req, forward this message to the single-threaded LAN only if it came from the LAN with the best master;
- when it receives from LAN A or LAN B a Delay_Req message, set the most significant two bits (MSB) of the sourcePortIdentity.portNumber of the Delay_Req message to “10” for LAN A and to “11” for LAN B respectively;
- when it receives from the single-threaded LAN a PTP message except Delay_Resp, forward a copy to both LAN A and LAN B;
- when it receives from the single threaded LAN a Delay_Resp, inspect its requestingPortIdentity.portNumber, restore it to the value it had in sourcePortIdentity.portNumber of Delay_Req and forward it only to the LAN where the Delay_Req came from.

A RedBox operating as a DATC when in P2P mode shall:

- forward a PTP message that it receives from LAN A or LAN B to the single-threaded LAN, but only if it came from the LAN with the best master and if it is neither a Delay_Req nor a Delay_Resp message;
- forward a copy of a PTP message that it receives from the single-threaded LAN to both LAN A and LAN B, except for Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req;
- send and receive the Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req messages and calculate the link delays as specified in IEC 61588:2009 on all ports individually.

A.4.5.5 PRP RedBox as stateless TC (SLTC)

A.4.5.5.1 PRP RedBox as SLTC – principle

In applications where the RedBox cannot implement the BMCA (e.g. simple hardware cores which forward frames in cut-through mode), the RedBox can be constructed as a stateless TC which does not discard PTP duplicates, but just sends all PTP messages it receives from LAN A or LAN B on the single-thread LAN C or LAN D, tagging them as coming from LAN A or LAN B by manipulating the bit 13 and bit 12 of the source port identity.

Singly attached clocks on the single-thread LAN C and LAN D will thus receive all PTP messages twice. Therefore, all OCs (e.g. OC6) execute the BMCA to pick up the correct Sync and Announce.

A.4.5.5.2 PRP RedBox as SLTC – limitations

This type of RedBox operates with E2E or P2P, 1-step or 2-step and over multiple domains. It expects that the OCs on the single-thread LAN C or LAN D follow the IEC 61588:2009 standard, and in particular the state machine in IEC 61588:2009, Figure 24, which prescribes that the slave clock only listens to the best master in case it receives Syncs from several masters.

This type of RedBox reserves the four most significant bits of the source port identity to tag the frames, therefore the number of ports per device is reduced to 4 096. This affects also clocks that are not redundancy-aware.

This method is vulnerable to a unidirectional link failure in the RedBox S to RedBox M direction because, if Delay_Req is lost Delay_Resp will not be received.

With this method, slave devices will report a portIdentity of the selected master which does not actually exist; this might cause some confusion to network administration.

An SLTC cannot be used to connect two duplicated HSR or PRP networks.

NOTE This method is therefore not applicable to slave-only clocks which do not execute the BMCA.

A.4.5.5.3 PRP RedBox as SLTC with E2E

Figure A.16 show the most complicated case in which the master OC5 sends an Announce or a Sync, which is duplicated by RedBox M and travels on LAN A and LAN B. When each of the two Announce or Syncs arrives at RedBox S (TC10), it tags them as coming from LAN A or LAN B by manipulating the source portNumber bit 13 and bit 12 and forwards them on LAN D.

OC6 receives two Announces and selects one of them as its best master. The different port identities serve as a tie-breaker. OC6 will from now on receive Syncs only from the selected master.

To measure link delay in end-to-end, OC6 sends a Delay_Req with its source port identity, which RedBox S receives. RedBox S manipulates the bit 13 and bit 12 of the source port identity when sending them over LAN A (value 10) and for LAN B (value 11).

NOTE 1 The reason for tagging the Delay_Req at RedBox S rather than at RedBox M as for a DATC is that an SLTC RedBox M cannot modify the tags since the SLTC RedBox S needs them.

Both Delay_Req travel to RedBox M, which sends them both unchanged to the master OC5. The master OC5 responds to each of them with a Delay_Resp.

NOTE 2 Master OC5 can notice that it receives redundant Delay_Resp since both come from a device with the same clockIdentity.


The RedBox M duplicates each Delay_Resp on both LANs, without modifying them. In a fault-free network four Delay_Resp are under way.

RedBox S can remove two of the Delay_Resp from the network, those where the tagging indicates that Delay_Req was on the other LAN (requestingPortIdentity not corresponding to the LAN the frame was received from), cases "10" and "01" in Figure A.17.

NOTE 3 If the RedBox cannot suppress the sending of the message, it can alternatively leave the requesting port identity untouched so the slave will not recognize this as a response to his request. Alternatively, it could also corrupt the CRC so the next network node will discard the frame.

RedBox S tags the Delay_Resp as it would tag the Syncs. Therefore, the OC will only take one and discard the other.


RedBox S resets the sets bit13 and bit 12 in the requesting port identity in the Delay_Resp back to "00" when sending the response to LAN D so the slave OC6 identifies this as his Delay_Resp to his Delay_Req. RedBox S does this only if the requesting port identity tagging in the Delay_Resp matches the source port identity tagging in the Delay_Req. Otherwise it leaves the requesting port identity untouched.



IEC

Figure A.16 – PRP RedBox as SLTC with E2E

The timing diagram in Figure A.17 shows the message exchange.



IEC

Figure A.17 – PRP RedBox as SLTC with E2E – timing

RedBox M can also operate as a DATC. In this case, it will send a Delay_Resp on the same LAN as it received the Delay_Req. It will then only send two Delay_Resp, one on each LAN A and LAN B, which will conserve the sourcePortIdentity.portNumber of Delay_Req in their requestingPortIdentity.PortNumber.

NOTE This was not necessary in a DATC, since a DATC does not operate in cut-through mode and can therefore select the messages it sends.

A.4.5.5.4 PRP RedBox as SLTC with P2P

The operation is similar to A.4.5.5.3, except that there are no Delay_Req and Delay_Resp messages (see Figure A.18). The peer delay messages are link-local and therefore unaffected. OC6 receives two Sync with a difference in the upper nibble and selects the best one as master.




Figure A.18 – PRP RedBox as SLTC with P2P

A.4.5.5.5 PRP RedBox as SLTC – specifications

A RedBox operating as SLTC shall operate as TC according to IEC 61588:2009.

A RedBox operating as SLTC, when operating in E2E, shall:

- if it receives a PTP message from a single-threaded LAN,
 - duplicate this message to LAN A and LAN B, and if the message is a Delay_Req, set the bit 13 and bit 12 in the most significant nibble (MSN) of the sourcePortIdentity.portNumber to “10” if it forwards it to LAN A and to “11” if it forwards it to LAN B, respectively.
- if it receives a PTP message from LAN A or LAN B,

- if it is a Delay_Resp message, discard it, if bit 13 and bit 12 of the requestingPortIdentity.portNumber do not match the value which the RedBox set in sourcePortIdentity.portNumber of Delay_Req on the LAN and if they match, reset these two bits to “00” before forwarding the Delay_Resp to the single-threaded LAN;
- otherwise, if it receives a PTP message from LAN A or LAN B in which the sourcePortIdentity bit 13 and bit 12 are 00, set them to “10” if it received it from LAN A, and to “11” if it received it from LAN B respectively, and forward the PTP message to the single-threaded LAN with possible correction.

A RedBox operating as SLTC, when operating in P2P, shall:

- forward a PTP message that it receives from LAN A or LAN B to the single-threaded LAN and if it is neither a Delay_Req nor a Delay_Resp message;
- forward a copy of a PTP message that it receives from the single-threaded LAN to both LAN A and LAN B, except for Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req;
- Send and receive the Pdelay_Resp (Pdelay_Resp_Follow_Up in 2-step) and Pdelay_Req messages and calculate the link delays as specified in IEC 61588:2009 on all ports individually.

A.5 HSR Mapping to PTP

A.5.1 PTP traffic in HSR

The operating mode of HSR is that a DANH sends the same frame in both directions into the ring, prefixed with the HSR Tag containing a sequence number. Each DANH relays the frame until it reaches the source again or is otherwise discarded. A destination DANH therefore receives two frames, uses the first and discards the duplicate.

This operating mode does not apply to PTP frames, since the delay suffered by a frame is not the same going clockwise or counter-clockwise in the ring. A DANH does not receive the same message from both ports since some messages are modified by intermediate nodes. Figure A.19 illustrates the paths of the different PTP messages in an HSR ring network.





Figure A.19 – HSR with one GMC


For 1-step operation, each DANH implements the 1-step Sync and P2P and exchanges the messages shown in Figure A.20.



IEC

Figure A.20 – PTP messages sent and received by an HSR node (1-step).

Alternatively, the HSR nodes may operate with 2-step correction as Figure A.21 shows.



IEC

Figure A.21 – PTP messages sent and received by an HSR node (2-step)

A.5.2 HSR nodes specifications

A DANH shall implement a Hybrid Clock, made of a TC and an optional OC.

A DANH shall send all PTP messages with a HSR tag.

NOTE 1 HSR tagging is in principle not needed for Pdelay_Req and Pdelay_Resp, which are link-local, but it allows to remove corrupted frames out of the ring, although in principle an HSR node does not forward untagged frames.

A DANH acting as master clock shall send all PTP messages in both directions with an HSR tag and with the same PTP sequence number.

A DANH shall handle all PTP messages for each port separately.

A DANH shall send and receive the Pdelay_Req and Pdelay_Resp messages per port and compute the link delay per port.

A DANH receives from each port one Announce and one Sync message coming from the same master with the same Clock Identity and different delays. It shall forward both in the same direction with the appropriate correction, except if this contradicts the HSR rules.

A clock shall apply the Best Master Clock algorithm on port A and port B, taking its ordinary clock as the common reference.

NOTE 1 TC-only DANH have been considered, but this would have added an unnecessary option.

If the Announce and Sync messages from both ports have the same Clock Identity, a DANP may use Syncs from both ports for synchronization/syntonization, or from one port only, a DANH shall listen to the port that provides the best clock quality.

A port that receives 2-step Sync message shall await the Follow_Up message from the same port before using the Sync timestamps to adjust its clock or to perform 2-step to 1-step conversion.

All PTP messages except the link-local Pdelay_Req/Pdelay_Resp shall use as source address the MAC address of the node that injected the frame into the ring. The link-local Pdelay_Req and Pdelay_Resp shall insert the MAC address of the TC.

NOTE 2 The IEEE P1588 revision group foresees that all devices that modify a frame insert their own MAC address as source address. Doing so would prevent removing circulating messages from the ring.

A.5.3 Redundant clocks in HSR

As Figure A.19 shows, there can be a redundant master clock in the ring. However, both cannot be active at the same time in PTP. Therefore, the grandmaster clock is selected among the master-capable clocks according to the Best Master Clock of PTP.

In case of loss of the master, the back-up master detects the loss and start broadcasting its own Announce messages.


A.5.4 Attachment of an MC to an external LAN

An MC may be connected to an HSR ring through a RedBox (see Figure A.22).

In its simplest form, this RedBox can be a three-port HSR node connected through an interlink port to an outside network.

The external network may operate either with a 1-step or with a 2-step clock.

The HSR node may be either slave or master. Such an HSR node may implement a TC or a BC between the interlink port and the ring. Figure A.22 shows that the RedBox at the same time translates from 2-step to 1-step clocks.



IEC

Figure A.22 – Attachment of a GMC to an HSR ring through a RedBox as TC

This arrangement allows to use an existing DANH to inject the synchronization. However it is not recommended when the MC is to be attached redundantly.

A.6 PRP to HSR Mapping

A.6.1 Connection methods

A PRP network is connected to an HSR ring through two RedBoxes. There are two different ways of proceeding for connecting the networks: through BCs (recommended) and through TCs (not recommended).

A.6.2 PRP-HSR connection by BC

The RedBoxes implement each a BC. The two BCs are treated as redundant clocks on HSR (see Figure A.23). The original PRP RCT is lost. Only one RedBox is actually sending Announces and Syncs, the other RedBox starts to send Announces and Syncs only in case it ceases to sense the other clock, but its port in the HSR ring remains otherwise in the PASSIVE_MASTER state (i.e. the second RedBox is not synchronized by the HSR ring).

Since both BCs reference the same GMC, they are only distinguished by their sourcePortIdentity.





Figure A.23 – PRP to HSR coupling by BCs

A.6.3 PRP-HSR connection by TCs

The redundant PRP-HSR connection by a TC causes the injection of four PTP messages into the ring, see Figure A.24. This configuration is not recommended and not further specified.




IEC

Figure A.24 – PRP to HSR coupling by TCs

A.7 Doubly attached clock model

A.7.1 State machine

Each port can be in the port states defined in IEC 61588:2009, Table 8, as Figure A.25 shows. The transitions that have been added to the state machines given in IEC 61588:2009, 9.2 for the Doubly Attached Clock are in red.



IEC

Figure A.25 – Port states including transitions for redundant operation

Table A.1 shows the states as defined in IEC 61588:2009, 9.2.5, Table 10 with additions (underscore) and deletions (strike-through).

Table A.1 – States

State	Definition
INITIALIZING	While a port is in the INITIALIZING state, the port initializes its data sets, hardware, and communication facilities. No port of the clock places any PTP messages on its communication path. If one port of a BC is in the INITIALIZING state, then all ports go to the INITIALIZING state.
FAULTY	The fault state of the protocol. A port in this state places no PTP messages except for management messages that are a required response to another management message on its communication path. In a BC, no activity on a faulty port affects the other ports of the device. If fault activity on a port in this state cannot be confined to the faulty port, then all ports go to the FAULTY state.
DISABLED	The port places no messages on its communication path. In a BC, no activity at the port affects the activity at any other port of the BC. A port in this state discards all PTP received messages except for management messages.
LISTENING	The port is waiting for the announceReceiptTimeout to expire or to receive an Announce message from a master. The purpose of this state is to allow orderly addition of clocks to a domain. A port in this state places no PTP messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, or signaling messages, or management messages that are a required response to another management message.
PRE_MASTER	The port behaves in all respects as though it were in the MASTER state except that it places no messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, signaling, or management messages.
MASTER	The port behaves as a master port.
PASSIVE_MASTER	The port places no messages on its communication path except for Pdelay_Req, Pdelay_Resp, Pdelay_Resp_Follow_Up, or signaling messages, or management messages that are a required response to another management message.
UNCALIBRATED	One or more master ports have been detected in the domain. The appropriate master port has been selected, and the local port is preparing to synchronize to the selected master port. This is a transient state to allow initialization of synchronization servos, updating of data sets when a new master port has been selected, and other implementation-specific activity.
SLAVE	The port is synchronizing to the selected master port.
PASSIVE SLAVE	The port is not synchronizing to the selected master port. In E2E, it may send Delay_Req; or in P2P, it sends Pdelay_Req, Pdelay_Resp, (Pdelay_Resp_Follow_Up). It may send signaling or management messages that are a required response to another management message.
NOTE “PASSIVE_MASTER” and “PASSIVE_SLAVE” are jointly expressed as “PASSIVE” in the management interface	

Table A.2 lists the transitions that cause state changes. In particular, the transition BMC_PSLAVE has been introduced to keep the paired ports synchronized.

Table A.2 – Transitions

Transition	Definition
ANNOUNCE_TO	No Announce has been detected on this port within the timeout. (the other port may still receive Announce messages)
BMC_SLAVE	Outcome of STATE_DECISION_EVENT: this port receives the best master. (the other port will go to PASSIVE_SLAVE)
BMC_MASTER	Outcome of STATE_DECISION_EVENT: the local OC is the best clock on both ports (the other port will also go to MASTER).
BMC_PASSIVE	Outcome of STATE_DECISION_EVENT: this OC is not the best but the port is prevented from entering the MASTER state (e.g. it is a grandmaster that cannot accept synchronization).
BMC_PSLAVE	Outcome of STATE_DECISION_EVENT: the clock on this port has a worse quality than that on the other port (the other port is in the SLAVE state).
DESIGNATED_DISABLED	A command has been given to disable this port.
DESIGNATED_ENABLED	A command has been given to enable this port.
FAULT_DETECTED	An error occurred. The logic is supposed to do some self-checking.
FAULT_CLEARED	The error has been corrected.
INITIALIZE	Reception of an INITIALIZE management message.
MASTER_CLOCK_SELECTED	The port has synchronized to the master on this port.
POWER_UP	The port starts operation.
QUALIFICATION_TO	The qualification time out (see IEC 61588:2009, 9.2.6.10) expired.
Ready	The initialization procedure is finished.
STATE_DECISION_EVENT	Outcome: BMC_MASTER, BMC_SLAVE, BMC_PSLAVE, BMC_PASSIVE
SYNC_FAULT	The clock synchronization has performed a correction outside of the allowed range.

Table A.3 indicates the auxiliary variables used in the state chart of Figure A.25.

Table A.3 – Variables

Variable	Definition
new_master	identity of the higher quality master
old_master	identity of the current master

A.7.2 Supervision of the port

A port in the SLAVE or PASSIVE_SLAVE state enters the FAULTY state when it ceases to receive the expected PTP messages. This governs the FAULTY_SYNC and the STATE_DECISION_EVENT transition.

To this effect, the ports shall implement a time-out on the PTP messages: Sync (and Follow_Up), Delay_Resp (in end-to-end) and Pdelay_Resp (in peer-to-peer).

The time-out on these messages shall be set by default to the double value of the expected period (logMessageInterval) plus 60 %.

NOTE 1 IEC 61588:2009 has no time-out on the PTP messages except Announce. This could change in a new edition.

NOTE 2 Only one outage is tolerated before the clock switches to the redundant path. If the Sync interval is 1 s, the time-out will be raised 2,6 s after reception of a correct Sync. This accounts for the ± 30 % jitter permitted by IEC 61588:2009, 7.7.2.1.

A.7.3 BMCA for paired ports

The BMCA algorithm in IEC 61588:2009 is extended for the purpose of paired ports synchronized by the same master. If BCs exist, the Steps Removed (number of hops) can be different and serve as a selection criteria.

The reception of identical messages is a normal case and therefore an application-specific quality chooses the clock to listen to, e.g. smaller correction field (after proper filtering) or an anomaly in the correction field.

The application-specific quality is reflected in the decision “*compare quality*” in Figure A.26, which has been adapted from Figure 28 of IEC 61588:2009.




Figure A.26 – BMCA for redundant masters

This paired port BMCA specification is considered in the term “better” transition in the sequel.

A.7.4 Selection of the port state

The paired port BMCA shall define which clock synchronizes which other clock.

- a) If the OC is the best (or only) master both on port A and port B, both ports go to the MASTER state (BMC_MASTER).
- b) If the master MCA on port A or the master MCB on port B is better than the OC, the port A with the better master goes to the SLAVE state (BMC_SLAVE) while the other goes to the PASSIVE_SLAVE state (BMC_PSLAVE).
- c) If a port in the SLAVE state loses its master (ANNOUNCE_TO, SYNC_FAULT), or the BMCA detects a better master on the paired port, the port transits from the SLAVE to the PASSIVE_SLAVE state (BMC_PSLAVE), while the other port wakes up from the PASSIVE_SLAVE state and becomes SLAVE (BMC_SLAVE).
- d) If the master on one port is worse than the OC or the master on the other port, no effort is spent in trying to supply the synchronisation, this is not a BC.

A.8 PTP datasets for high availability

A.8.1 General

The objects in the IEC 61588 datasets are extended. These objects are exposed in the MIB for local or remote management. In case of discrepancy, the wording in the MIB takes precedence.

A.8.2 Data types

A.8.2.1.1 TimeInterval

This type is defined in IEC 61588, 5.3.2

A.8.2.1.2 EUI64

Base type: OCTET STRING (SIZE(8))

Description: This type is defined in IEC 61588, 7.5.2.2.2

A.8.2.1.3 PortStateType

This type is defined in IEC 61588:2009, 8.2.5.3.1, Table 8 and extended by value 10: PASSIVE_SLAVE.

A.8.2.1.4 ProfileSet

Base type: Unsigned32 (bit set)

References: IEC 62439-3:2015

Description: Bits are numbered from 0 (least significant) to 31 (most significant). Bit “7” is set for the L3E2E profile (see Clause C.5), Bit “8” for the L2P2P profile (see Clause C.6) and Bit “9” for the IEC/IEEE 61850-9-3 profile (see Annex B). Other bits are not defined. If the value of this object is “0”, no profile is defined.

A.8.2.1.5 PortAttachmentType

Base type: Enum32

Values:

- 0: not specified;
- 1: the port belongs to a singly attached ordinary clock;
- 2: the port belongs to a singly attached boundary clock;
- 3: the port belongs to a singly attached transparent clock;
- 4: the port belongs to a boundary node;
- 5: the port belongs to a doubly-attached ordinary clock (portDs.pairedPort is not >FFFF);
- 6: the port belongs to a doubly-attached boundary clock (portDs.pairedPort is not >FFFF);
- 7: the port belongs to a doubly-attached transparent clock (portDs.pairedPort is not >FFFF);
- 8: the port belongs to a stateless transparent clock (portDs.pairedPort is not >FFFF);
- 255: unknown.

A.8.3 Datasets for ordinary or boundary clocks

A.8.3.1 Additional or modified members

The members additional or modified with respect to IEC 61588 are as follows:

- defaultDs.numberPorts
- defaultDs.profileSet;
- defaultDs.timeInaccuracy;
- portDs.portState;
- portDs.portEnabled;
- portDs.dlyAssymetry;
- defaultDs.profileId;
- portDs.vlanEnable;
- portDs.vlanId;
- portDs.vlanPrio;
- portDs.twoStepFlag;
- portDs.peerIdentity;
- portDs.prpPairedPort;
- portDs.prpAttachment;
- portDs.errorCounter;
- portDs.peerDelayMax.

A.8.3.2 Dataset DefaultDS

A.8.3.2.1 Object defaultDs.numberPorts

Type: Unsigned 8 [5 to 254]

Reference: IEC 61588:2009, 8.2.1.2.3, IEC 62439-3:2015, Clause A.3

Description: The number of PTP ports on the device. (1 if OC, >1 for Boundary Clock). IEC 62439-3 allows 2 for OC that are doubly attached.

A.8.3.2.2 Object defaultDs.profileSet

Type: ProfileSet

Description: The set of profiles which the clock has the ability to support with the required performance. Profile activation is a port property.

Default value: 0

NOTE This object serves management to check that all devices are properly configured for deployment or maintenance. This information is not used in the PTP protocol. It does not guarantee that all parameters of the clock have been properly set. A clock can support several profiles, but a port supports only one.

A.8.3.2.3 Object defaultDs.timeInaccuracy

Type: TimeInterval

Units: scaledNanoseconds

Reference: IEC 61588:2009, 5.3.2, IEC 62439-3:2015, C.7.4.1 and C.7.6.1.

Description: The absolute value of the rated worst-case time error with respect to absolute time when the clock is connected to its reference signal. If the clock is a boundary clock, the reference signal is the clock signal applied to any port, and the rated time error the error observed on any master port.

Default value: 0xFFFFFFFF

NOTE This object serves to check that a MC or BC has a sufficient quality for the given application. This value is not measured at run-time, it is a design parameter that also appears in the MIB. This value is not transmitted using in the PTP protocol. The inaccuracy at run-time can be deduced from the defaultDS.clockQuality.offsetScaledLogVariance object.

A.8.3.2.4 Object defaultDs.offsetFromMasterLim

Type: Unsigned32

Units: nanoseconds

Reference: IEC 62439-3:2015, A.8.3.2.4

Description: The value of the limit to the offset from master that signals an abnormal condition.

Default value: 0xFFFFFFFF

A.8.3.3 Dataset ParentDS

As in IEC 61588.

A.8.3.4 Dataset PortDS

A.8.3.4.1 Object portDs.portState

Type: PortStateType

Reference: IEC 61588:2009, 8.2.5.3.1, Table 8; IEC 62439-3:2015, Table A.1

Description: The current state of the protocol engine associated with this port taken from the enumeration in Table 8, extended by the PASSIVE_SLAVE state in doubly attached clocks.

A.8.3.4.2 Object portDs.portEnabled

Type: Boolean

Reference: IEC 61588:2009, 8.2.5.3.1

Description: Enables the port for PTP traffic and indicates the enabled status of the port.

Default value: true

A.8.3.4.3 Object portDs.dlyAsymmetry

Type: TimeInterval

Units: scaledNanoseconds – IEC 61588, 5.3.2

Reference: IEC 61588:2009, 7.4.2, 11.6, 5.3.2

Description: The link delay asymmetry to correct the link delay.

Default value: 0

NOTE This object is used to verify the topology in P2P.

A.8.3.4.4 Object portDs.profileId

a) Type: ProfileSet

b) Reference: IEC 62439-3:2015

c) Description: The profile in use on this port

Default value: 0

NOTE This object serves to check that all devices are properly configured for deployment or maintenance. This information is not used in the PTP protocol. It does not guarantee that all parameters of the clock have been properly set.

A.8.3.4.5 Object portDs.vlanEnable

Type: Boolean

Reference: IEEE 802.1Q:2014

Description: Enables that PTP messages generated internally to the clock receive an IEEE 802.1Q tag with the values specified in portDs.vlanId and portDs.vlanPrio when sent over this port, with exception of Pdelay_Req and Pdelay_Resp, which operate with link-local addresses. It does not apply to messages in transit generated by other clocks. The bridge rules take precedence over this setting.

Default value: false

NOTE The VLAN settings of the port in its bridge function can override this tagging or remove it. A bridge port cannot send a VID of 0 except if its VLAN setting is disabled.

A.8.3.4.6 Object portDs.vlanId

Type: Unsigned32 (0 to 4095)

Reference: IEEE 802.1Q:2014

Description: The IEEE 802.1Q VLAN Identity to be used by the PTP stack internal to a bridge if portDs.vlanEnable is true.

Default value: 0

NOTE A bridge port would not send a message with a VID of 0 if the port VLAN setting indicates that VLANs are enabled.

A.8.3.4.7 Object portDs.vlanPrio

Type: Unsigned32 (0 to 7)

Reference: IEEE 802.1Q:2104

Definition: The IEEE 802.1Q priority (0 to 7) to be used by PTP messages originating from the PTP stack of this clock if portDs.vlanEnable is true.

Default value: 4

A.8.3.4.8 Object portDs.twoStepFlag

Type: Boolean

Reference: IEC 61588:2009, 8.2.1.2.1

Description: Enables (if true) 2-step correction at sending of this port (if the port implements it) and indicates that this port uses 2-step correction.

Default value: true

NOTE A BC can use a different correction method on different ports, according to the supported profile.

A.8.3.4.9 Object portDs.peerIdentity

Type: OCTET STRING (SIZE (6))

Reference: ISO/IEC/IEEE 8802-3:2014

Description: The Medium Access Control address of the peer clock in a peer-to-peer delay measurement.

A.8.3.4.10 Object portDs.prpAttachment

Type: PortAttachmentType

Reference: IEC 62439-3:2015

Description: The type of redundant attachment of this port.

Default value: 0

A.8.3.4.11 Object portDs.prpPairedPort

Type: Unsigned16 [0 to 4095]

Description: The portNumber of the other port to which this port is attached, if the port is paired for the purpose of double attachment; a value greater than 4 096 indicates the port is not paired.

Default value: 65535

A.8.3.4.12 Object portDs.errorCounter

Type: Counter32

Description: The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that are normally received given the port state.

Default value: 0

A.8.3.4.13 Object portDs.peerDelayLim

Type: PtpTimeInterval

Description: The value of the peer delay that, if exceeded, will cause an event to be raised.

Default value: 100 ns

A.8.4 Object for transparent clocks

A.8.4.1 Additional or modified objects

The members additional or modified with respect to IEC 61588 are as follows:

- transparentClockDefaultDS.profileSet;
- transparentClockDefaultDS.timeInaccuracy
- transparentClockPortDs.portEnabled;
- transparentClockPortDs.dlyAssymetry;transparentClockPortDS.twoStepFlag;
- transparentClockPortDs.peerIdentity;
- transparentClockPortDs.prpAttachment;
- transparentClockPortDs.prpPairedPort;transparentClockPortDs.errorCounter;
- transparentClockPortDs.delayMax.

A.8.4.2 Dataset TransparentClockDefaultDS

A.8.4.2.1 Object transparentClockDefaultDS.profileSet

This object has the same definition as in A.8.3.2.2.

A.8.4.2.2 Object transparentClockDefaultDS.timeInaccuracy

Type: TimeInterval

Units: scaledNanoseconds – IEC 61588:2009, 5.3.2

Reference: IEC 62439-3:2015.

Description: The absolute value of the rated worst-case time error introduced between any two ports.

Default value: 50 ns

NOTE This object serves to check that a TC has a sufficient quality for the given application. This value is not measured at run-time, it is a design parameter that also appears in the MIB. This value is not transmitted using in the PTP protocol.

A.8.4.3 Dataset transparentClockPortDs**A.8.4.3.1 Object transparentClockPortDs.portEnabled**

Type: Boolean

Reference: IEC 61588:2009, 8.2.5.3.1

Description: Enables the port for PTP traffic and indicates the enabled status of the port.

Default value: true

A.8.4.3.2 Object transparentClockPortDs.dlyAsymmetry

This object has the same definition as in A.8.3.4.2.

A.8.4.3.3 Object transparentClockPortDs.twoStepFlag

Type: Boolean

Reference: IEC 61588:2009, 8.2.1.2.1

Description: Enables (if true) 2-step correction at sending of this port (if the port implements it) and indicates that this port uses 2-step correction.

Default value: true

NOTE A BC can use a different correction method on different ports, according to the supported profile.

A.8.4.3.4 Object transparentClockPortDs.peerIdentity

Type: OCTET STRING (SIZE (6))

Reference: ISO/IEC/IEEE 8802-3:2014

Description: The Medium Access Control address of the peer clock in a peer-to-peer delay measurement.

A.8.4.3.5 Object transparentClockPortDs.prpAttachment

Type: PortAttachmentType

Reference: IEC 62439-3:2015

Description: The type of redundant attachment of this port.

Default value: 0

A.8.4.3.6 Object transparentClockPortDs.prpPairedPort

Type: Unsigned16 [0 to 4095]

Description: The portNumber of the other port to which this port is attached, if the port is paired for the purpose of double attachment, a value greater than 4 096 indicates the port is not paired.

Default value: 65 535

A.8.4.3.7 Object transparentClockPortDs.errorCounter

Type: Counter32

Description: The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that are normally received given the port state.

Default value: 0

A.8.4.3.8 Object transparentClockPortDs.peerDelayLim

Type: PtpTimeInterval

Description: The value of the peer delay that, if exceeded, will cause an event to be raised.

Default value: 100 ns

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

Annex B (normative)

PTP profile for Power Utility Automation – Redundant clock attachment

B.1 Application domain

For Power Utility Automation, IEC/IEEE 61850-9-3 specifies a Precision Time Protocol profile using a layer 2 transport according to IEC 61588:2009, Annex F with a peer-to-peer delay measurement based on IEC 61588:2009, Clause J.4.

Annex B defines the mandatory support of redundant clock attachment as an extension of IEC/IEEE 61850-9-3:—, Clause 9.

B.2 PTP profile specification

The profile identifier of a singly or doubly attached clock according to IEC/IEEE 61850-9-3 and Annex A shall be:

profileName: IEC/IEEE 61850-9-3 "Precision Time Protocol profile for Power Utility Automation"

profileVersion: 1.0

profileIdentifier: 00-0C-CD-00-01-xy

Where the first nibble of the sixth octet is a bitfield:

- x = 0 for a singly-attached clock
- x = 1 for PRP redundancy
- x = 2 for HSR redundancy
- x = 3 for both PRP and HSR (configurable) redundancy
- y = 0 (minor revision)

organizationName: IEC Technical Committee 57 Working Group 10

sourceIdentification: A copy of IEC/IEEE 61850-9-3 can be obtained from the International Electrotechnical Commission (<http://www.iec.ch>)

B.3 Redundant clock attachment

Clocks shall implement the doubly attached clock model for at least one of:

- a) PRP (A.4.2) or
- b) HSR (A.5) or
- c) HSR-PRP RedBoxes (A 6)

The manufacturer shall list the supported method in the PICS, as stated in IEC/IEEE 61850-9-3.

Annex C (normative)

PTP profiles for high-availability automation networks

C.1 Application domain

For high-availability industrial automation networks, two precision time clock profiles are specified which share most specifications:

- L3E2E, transported over Layer 3 using end-to-end delay measurement (E2E), based on IEC 61588:2009, Clause J.3 and
- L2P2P, transported over Layer 2 using peer-to-peer delay measurement (P2P), based on IEC 61588:2009, Clause J.4.

An objective is that clock synchronization operates without special configuration if all network elements comply with the same profile.

Interoperability between L3E2E and L2P2P clocks is not an objective, but coupling them through BCs is possible. Clocks of both profiles can coexist on the same network.

These profiles allow mixing singly attached clocks according to the IEC 61588:2009 default profiles and doubly-attached clocks according to Annex A in the same network.

C.2 PTP profile specification

The profile specification of a singly or doubly attached clock according to Annex C shall be according to IEC 61588:2009, 19.3.3:

profileName: IEC 62439-3 "Precision Time Protocol profile for High-availability Automation Networks"

profileVersion: 1.0

profileIdentifier: 00-15-4E-00-01-xy
 Where the first nibble of the sixth octet is a bit field:
 x = 1 for PRP redundancy
 x = 2 for HSR redundancy
 x = 4 for L3E2E
 x = 8 for L2P2P
 y = 0 (minor version)

EXAMPLE: 00-15-4E-00-01-50 supports L3E2E and PRP.

organizationName: IEC Technical Committee SC65C Working Group 15

sourceIdentification: A copy of IEC 62439-3 can be obtained from the International Electrotechnical Commission (<http://www.iec.ch>)

NOTE The OUI defined in IEC 61588:2009, Clause J.4 does not apply when redundancy is used.

C.3 Clock types

This specification distinguishes clocks by their capabilities:

- Ordinary Clocks (OCs) implement one of the following capabilities:
 - Slave-only (defaultDS.slaveOnly = true)
no port can be in the MASTER state

- Grandmaster-capable (defaultDS.slaveOnly = False); the port (or the port pair in redundancy) can be in the MASTER state
- Grandmaster-only (defaultDS.slaveOnly = False and clockClass = 6 or 7) no port can be in the SLAVE state.
- Transparent Clocks (TCs) correct and forward PTP messages; ports do not have states in IEC 61588:2009.
- Boundary Clocks (BC) are never slave-only and can have either:
 - exactly one port in the SLAVE state and all other in the MASTER or PASSIVE state;
 - all ports in the MASTER state, in which case the BC is the grandmaster.

NOTE Within a time domain, a BC has a port in the SLAVE state in the upper sub-domain and one or several ports in the MASTER state in the lower sub-domain(s) (see IEC TR 61850-90-4).

C.4 Protocol specification common

All clocks shall support the PTP time scale (based on TAI).

All singly attached clocks shall support the default best master clock algorithm in IEC 61588:2009, 9.3.2, 9.3.3 and 9.3.4.

All doubly attached clocks according to IEC 62439-3 shall support in addition the extension to the best master clock algorithm defined in A.7.3.

All clocks shall support both 1-step and 2-step synchronization on ingress, they may use 1-step or 2-step synchronization on egress.

All clocks shall support at least one of the three management mechanisms:

- 1) the alternate management mechanism using the SNMP MIB specified in Annex E and/or
- 2) the alternate management mechanism using the management objects defined in IEC TR 61850-90-4:2013, 19.3 and 19.4 and/or
- 3) the manufacturer-defined fixed values and/or the manufacturer-specific implementation method to access all configurable values as stated in IEC 61588:2009, 15.1.1.

C.5 Protocol specification for L3E2E automation profile

All clocks shall transport PTP messages using Layer 3 (UDP over IP) according to IEC 61588:2009, Annex D (IPv4) or IEC 61588:2009, Annex E (IPv6)

All clocks shall use only multicast addresses, as specified in IEC 61588:2009, Clause D.3 and IEC 61588:2009, Clause E.3 respectively.

All clocks shall support the end-to-end delay calculation defined in IEC 61588:2009, 10.2.

The “portnumber” of all clocks shall be zero in their four most significant bits.

C.6 Protocol specification for L2P2P automation profile

All clocks shall transmit PTP messages on Ethernet layer 2 according to IEC 61588:2009, Annex F.

All clocks shall use multicast communication using the addresses specified in IEC 61588:2009, Clause F.1.

All clocks shall support the peer-to-peer delay measurement defined in IEC 61588:2009, 10.3 and 11.4.

C.7 Timing requirements

C.7.1 Measurement conditions

Steady state is defined as 30 s after a single master starts to send synchronisation messages and 16 s after a change of master, with no change to the environment temperature.

This requirement applies only to clocks that have been energized for 30 min to accommodate for instance temperature-controlled oscillators.

NOTE Assuming a Gaussian distribution, this corresponds to three sigma ($3\sigma = 99.7\%$) or no more than 3 points outside the specified interval, out of 1 000 total points evaluated.

C.7.2 Network time inaccuracy

The following requirements on the network elements aim at achieving a network time inaccuracy better than $\pm 1\ \mu\text{s}$ after crossing approximately 15 TCs or 3 BCs.

C.7.3 Network elements

All network elements shall be clocks according to this standard, conformant media converters and/or conformant links.

C.7.4 Requirements for grandmasters

C.7.4.1 Grandmaster time inaccuracy

A grandmaster-capable clock shall have a time inaccuracy measured between its applied time reference signal and the produced synchronization messages that is smaller than 250 ns.

NOTE This value corresponds to an IEC 61588:2009 clockAccuracy of 22 hex; IEC 61588 defines clockAccuracy.

In case the grandmaster-capable clock has no time reference signal, IEC 61588:2009, J.4.4.1 shall apply.

C.7.4.2 Grandmaster holdover

A grandmaster shall remain within the time inaccuracy of C.7.4.1 for a holdover time of at least 5 s after losing its time reference signal, given that it was in steady state.

C.7.4.3 Grandmaster clockClass in holdover and recovery

A grandmaster clock shall have an IEC 61588:2009 clockClass of:

- 6 while connected to its time reference signal and in steady state;
- 7 after loss of time reference signal, while in holdover;
- 52 when its time error exceeds the value given in C.7.4.1;
- 187 when its time error exceeds $1\ \mu\text{s}$;
- 6 after recovering the time reference signal and in steady state.

NOTE 1 This modifies IEC 61588:2009, Table 5 for the timing requirements of this profile.

NOTE 2 The timeTraceable bit of the flagField indicates if the grandmaster is currently locked to a recognized standard time source or has a stated uncertainty to a recognized standard time source. The clockClass = 6 indicates that the grandmaster is currently locked to its time reference signal. The clock class 6 appears twice, once before loss of time reference signal, and once after recovery of that signal.

C.7.5 Requirements for TCs

A TC shall introduce less than 50 ns of time inaccuracy, measured between the applied synchronization messages at any ingress port and the produced synchronization messages at any egress port, given it is in steady state.

NOTE This time inaccuracy cumulates the time errors in measuring the residence delay, measuring the peer delay in the ingress port and responding to the peer delay measurement from the downstream clocks in the egress port(s). It does not include time error due to asymmetry in either link, nor introduced by the upstream or downstream peer devices involved in the peer delay measurements. This time inaccuracy can be measured using peer devices with known contribution of the peer delay measurement to time error and using links with known or negligible asymmetry, and subtracting those amounts from the observed TC time error.

C.7.6 Requirements for BCs

C.7.6.1 BC time inaccuracy

A BC shall introduce less than 200 ns of time inaccuracy between the port in the SLAVE state and any port in the MASTER state, provided that it is in steady state.

NOTE This time inaccuracy cumulates the time errors in local clock adjustment, measuring the peer delay in the port in the SLAVE state and responding to the peer delay measurement from the downstream clocks in the port(s) in the MASTER state. It does not include time errors due to asymmetry in either link, nor introduced by the upstream or downstream peer devices involved in the peer delay measurements. This time inaccuracy can be measured using peer devices with known contribution of the peer delay measurement to time error and using links with known or negligible asymmetry, and subtracting those amounts from the observed BC time error.

C.7.6.2 BC as free-running grandmaster

In case the BC has no port in the SLAVE state and no time reference signal, IEC 61588:2009, J.4.4.1 shall apply.

C.7.6.3 BC as master in holdover

A BC shall remain within the time inaccuracy of C.7.6.1 for a holdover time of at least 5 s after losing its time reference signal or PTP synchronization, provided that it was in steady state.

C.7.7 Requirements for media converters

Media converters that support IEC 61588 are considered as TCs or BCs and are subject to the requirements of C.7.5 and C.7.6, respectively.

Media converters (e.g. fibre to copper) introduce a delay, considered in peer-to-peer delay measurement. This delay may suffer from a significant jitter and may be different in both directions. Network engineering should use only media converters with known behaviour to estimate the network time inaccuracy at every slave clock.

Media converters shall present a jitter of less than 50 ns and an asymmetry of less than 25 ns.

C.7.8 Requirements for links

Links present a predictable and nearly constant link propagation delay (about 5 μ s/km for fibre or copper), whose average value is regularly calculated by peer-to-peer delay measurement.

NOTE Radio links are not considered.

Since delay asymmetry is not measurable, network engineering must know it to compensate it.

Links shall present a propagation asymmetry of less than 25 ns or shall have a known propagation asymmetry with an asymmetry variation of less than 25 ns.

C.8 Network engineering

To achieve the required network time inaccuracy, careful network design is required, considering the placement of masters and redundant masters and possible network topology changes because of reconfigurations so as not to exceed the number of allowed TCs and BCs.

The network designer should only select network elements knowing their contribution to network time inaccuracy and dependencies on the operating conditions. The network designer should estimate the network time inaccuracy for all slave clocks.

The network designer should use network elements with stricter specifications in more demanding applications or larger networks.

The total time inaccuracy ε_T of the time signal available at a given clock is computed at engineering time as:

$$\varepsilon_T = \varepsilon_{RF} + \varepsilon_{GM} + N_{TC} \times \varepsilon_{TC} + N_{BC} \times \varepsilon_{BC} + N_{MC} \times \varepsilon_{MC}$$

The network time inaccuracy ε_N of the time signal available at a given clock is computed at engineering time as:

$$\varepsilon_N = \varepsilon_{GM} + N_{TC} \times \varepsilon_{TC} + N_{BC} \times \varepsilon_{BC} + N_{MC} \times \varepsilon_{MC}$$

where:

ε_T = total time inaccuracy;

ε_N = network time inaccuracy;

ε_{RF} = worst-case source time inaccuracy;

ε_{GM} = worst-case grandmaster time inaccuracy (see C.7.4);

ε_{TC} = worst-case TC time inaccuracy (see C.7.5);

ε_{BC} = worst-case BC time inaccuracy (see C.7.6);

ε_{MC} = worst-case media converter time inaccuracy (see C.7.7);

N_{TC} = number of TCs in series on the longest path to this clock;

N_{BC} = number of BCs in the path on the longest path to this clock;

N_{MC} = number of media converters in the path on the longest path to this clock.

The network engineering should consider the time error ε_{RF} in the signal received from the recognized standard time source, e.g. the GNSS signal.

The network commissioner should check the time inaccuracy class of the installed network components and verify the topology (e.g. using the management mechanism, if such is provided) and then perform a calibration for the network time inaccuracy of all slave clocks. For this purpose, it is recommended to equip clocks with a 1 PPS output.

C.9 Default settings

All clocks shall be initially configured with the default settings and shall be able to support the ranges (if specified) of Table C.1.

Table C.1 – PTP attributes for the Industrial Automation profile

PTP attribute	Default value	Range
defaultDS.domainNumber	0	as specified in Table 2 of IEC 61588:2009
portDS.logAnnounceInterval	1 for the L3E2E profile 0 for the L2P2P profile	0 to 4 for the L3E2E profile 0 for the L2PTP profile ^a
portDS.logSyncInterval	0	0 ^a
portDS.logMinPdelay_ReqInterval	0	0 ^a
portDS.announceReceiptTimeout	3	3
defaultDS.priority1	255 for slave-only clocks, 128 otherwise	0 to 255 255 for slave-only clocks
defaultDS.priority2	255 for slave-only clocks, 128 otherwise	0 to 255 255 for slave-only clocks
defaultDS.slaveOnly	False (if this parameter is configurable)	True, False
transparentClockdefaultDS.primaryDomain	0	as specified in Table 2 of IEC 61588:2009
τ (Allan deviation sample period)	1,0 s	1,0 s
For each defined range, manufacturers are free to allow wider ranges, which are out of the scope of this profile and which conformance testing will not consider.		
^a Although a master clock is not required to support any other value, it still can have a jitter as IEC 61588:2009, 9.5.9.2 specifies. A slave clock should have a capture range accordingly.		

C.10 Redundant clock handling

Clause C.10 does not apply to clocks with only one port.

These profiles differ from IEC 61588:2009 when clocks are doubly attached according to Annex A as follows:

- 1) Introduction of ordinary clocks with more than one port;
- 2) Extension of the BMCA to allow a decision based on an application-specific clock quality when the same grandmaster appears over two different paths;
- 3) Extension of the port state machine for doubly attached clock to allow checking of unused paths;
- 4) Time-outs for supervision of all message exchanges;
- 5) Extension of the datasets to include redundancy management.

When redundancy is used, it shall be implemented as follows:

Clocks shall implement the doubly attached clock model for at least one of:

- a) PRP (A.4.2) or
- b) HSR (A.5) or
- c) HSR-PRP RedBoxes (A 6)

The manufacturer shall list the supported method in the PICS.

C.11 Protocol Implementation Conformance Statement (PICS)

C.11.1 Conventions

The conventions are based on ISO/IEC 9646-7.

The "Base" column reflects the definitions and specifications in the base standard (i.e. Clauses C.3 to C.10 of this document in this case). Each entry in this column is chosen from the following list:

- **mandatory (m):** the base standard mandates this capability and it is implemented;
- **optional (o):** the base standard leaves this capability optional, but it is implemented.

The "Condition" column reflects the condition(s) for the capability. Each entry in this column is chosen from the following list:

- **conditionally supported (c):** this capability is supported under the conditions specified in the corresponding index in the PICS Table;
- **unconditionally supported (–):** this capability is not constrained.

The "Support" column is to be filled by the manufacturer.

C.11.2 PICS

The manufacturer shall list the capabilities and options supported by its clock in the PICS according to Table C.2.

The "Base" of this PICS are Clauses C.3 to C.10 of this document.

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016

Table C.2 – PICS for clocks

PICS proforma reference	Capability	Value range	Base	Condition ^a	Support (see Note)
CLOCK_TYPE_OC	clock is OC according to this base	True, False	m	c.1	<i>True</i>
CLOCK_TYPE_TC	clock is TC according to this base	True, False	m	c.1	
CLOCK_TYPE_BC	clock is BC according to this base	True, False	m	c.1	
NR_PORTS	number of clock ports (total)	integer > 0	m	-	2
PORTS_STEP	1: all ports support 1-step on egress 2: all ports support 2-step on egress 3: all ports support both 1-step and 2.	1 to 3	m	-	{1}
SLAVE_ONLY	all ports of the clock are slave-only	True, False	m	c.2	<i>True</i>
TIME_TRACEABLE	connectable to a time reference outside of PTP (e.g. GPS)	True, False	m	c.3	
FREQ_TRACEABLE	connectable to a frequency reference outside of PTP (e.g. GPS)	True, False	m	c.3	
DAC	doubly attached OC	True, False	o	-	<i>True</i>
PORTS_PAISED	paired clock ports for redundancy (e.g. {3-4})	identifier pair	o	c.4	{0, 1}
REDBOX_DATC	Redbox as TC	True, False	o	c.5	
REDBOX_SLTC	Redbox as Stateless TC	True, False	o	c.5	
REDBOX_TWBC	Redbox as three-way BC	True, False	o	c.6	
REDBOX_DABC	Redbox as DAC BC	True, False	o	c.6	
MIB_SNMP	supports MIB of IEC 62439-3, Annex E	True, False	m	c.7	
MIB_61850	supports IEC TR 61850-90-4 Clock Objects	True, False	m	c.7	
MIB_OTHER	clock supports fixed values or a mechanism defined by the manufacturer (If True, this list is appended to this PICS)	True, False	m	c.7	<i>True, void list.</i>
NOTE The values in the “Support” column are an example of a clock that is doubly attached over ports 0 and 1, uses 1-step PdelayReq, but has no master capability and does not support an MIB (there is nothing adjustable). ^a Conditions are: c.1: at least one shall be supported (CLOCK_TYPE_OC and CLOCK_TYPE_TC may be both True) c.2: only if CLOCK_TYPE_OC = True c.3: only if SLAVE_ONLY = False c.4: shall be “m” (>1) if DAC = True c.5: support shall only be declared if CLOCK_TYPE_TC = True and DAC=True c.6: support shall only be declared if CLOCK_TYPE_BC = True and DAC=True c.7: at least one shall be supported					

Clock accuracy is the maximum guaranteed error from the reference time (in Figure D.1: 200 ns)

Quality criteria for non-synchronized clocks are the stability, expressed for instance in $\mu\text{s/s}$, and the temperature stability ($\mu\text{s/s/}^{\circ}\text{C}$).

A clock that is only synchronized to a reference frequency is called a syntonized clock.

NOTE Similar considerations apply to frequency, which is not considered here.

D.3 PTP clock types

The basic principle of IEC 61588:2009 is shown in Figure D.2.




Figure D.2 – Precision Time Protocol principle

A grandmaster clock (GMC) distributes the exact time in a Sync message multicast to all (slave) ordinary clocks (OC). This Sync message suffers different delays on the path between the grandmaster and the slaves, especially a residence delay when traversing the bridges and a link delay on the medium.


The bridges operate here as TCs capable of calculating the residence delay and correcting the time for them. To this effect, each TC (and each master and each slave) has hardware capable of time-stamping the PTP messages that transit through them with nanosecond precision.

More generally, PTP distinguishes the following clock types (see Figure D.3):

- Ordinary clocks (OC) have only one port (singly-attached clock), that operates either in the slave or in the master state. An ordinary clock can be configured as master-capable or

slave-only. When the OC is the best clock of its subdomain, it becomes the Grandmaster Clock (GMC); its port is in the MASTER state and sends Sync messages to synchronize all slave clocks. Otherwise, its port is in the SLAVE (or PASSIVE) state and it receives synchronization messages.

- Transparent clocks (TC) have at least two ports which are stateless, and acts as a bridge for all messages. A TC forwards the Sync messages it receives over one port to all its other ports and corrects the time indicated in the Sync messages.
- Boundary clocks (BC) have at least two ports, that are usually in different states, exactly one port being in the SLAVE state and synchronized by a master clock while the other port(s) is (are) in the MASTER state and synchronize(s) a subdomain. When the BC receives a Sync message over the port in the SLAVE state, it does not forward it to the ports in the MASTER state; instead, it sends its own PTP messages, with a period that may be different from port to port. If the BC is the best clock of all subdomains, all its ports are in the MASTER state and it becomes the Grandmaster clock. Otherwise, it becomes master in each subdomain where it is the best master. It may be degraded to SLAVE or PASSIVE by a better master in its subdomain.
- Hybrid clocks (HC) combine a transparent clock (TC) and an ordinary clock (OC).



IEC

Key:

The text in **red** expresses the momentary role of a clock

The text in **blue** expresses the current state of a port according to its state machine.

The text in **black** is the configured or factory type of clock

Figure D.3 – Precision Time Protocol elements

D.4 PTP main options

IEC 61588:2009 distinguishes three basic options:

- layer 2 versus layer 3 communication (see D.5)
- 1-step versus 2-step clock correction (see D.5).
- end-to-end (see D.7) versus peer-to-peer link delay measurement (see D.8)

D.5 Layer 2 and layer 3 communication

Layer 2 PTP communication transports PTP message on layer 2, relying on MAC addresses. This communication ends at a router. Layer 2 PTP defines for Ethernet at which point in a frame the time stamps is taken. Layer 2 PTP is the preferred approach for automation networks based on Ethernet Layer 2.


Layer 3 PTP communication transports PTP messages over layer 3, relying on IPv4 or IPv6 addresses (mapped to specific Layer 2 addresses). Layer 3 PTP can operate over a variety of media. The precise time-stamping relies on an unspecified physical layer support. Layer 3 PTP is the preferred approach for automation networks based on Layer3.

D.6 1-step and 2-step correction

D.6.1 Time correction in TCs

The master clock (MC) multicasts periodically (e.g. every 1 s) a Sync message containing the reference time. This Sync message transits through the network over links and TCs.

The Sync message suffers a path delay consisting of the link delays λ between clocks and of the residence delays ρ in the TCs, before it reaches the slave clocks. To correct these delays, the network devices time-stamp precisely the Sync messages as they ingress and egress the devices (in the physical interface, PHY) as shown in Figure D.4.



IEC

Figure D.4 – Delays and time-stamping logic in TCs

The TCs calculate a correction including the residence time of the Sync message. There are two ways to transmit this correction: 1-step correction and 2-step correction, which are illustrated in Figure D.5.

- 1) 1-step correction: the TC enters the value of the correction in the Sync message while it is being sent. This needs a hardware support that can modify messages on-the-fly.
- 2) 2-step correction: the TC sends the value of the correction in a subsequent Follow_Up message. The hardware support is limited to the precise time-stamping of the ingress and egress frames.




Figure D.5 – Correction of the Sync message by 1-step and 2-step (peer-to-peer)

Both methods can be used within a network (see D.6.2).

D.6.2 2-step to 1-step translation

D.6.2.1 Choice of methods

IEC 61588:2009 allows to mix clocks with 1-step correction and 2-step correction within the same subdomain.

The engineer is faced with two correction methods, depending on the ability of its devices. Most current commercial clocks use 2-step correction to simplify implementation, but with increasing silicon support 1-step is gaining ground.

In HSR rings, 1-step is preferred since hardware support is already provided for cut-through and handling Follow_Ups would only be overhead.

Clocks supporting the two correction methods can be chained in the same network if they accept both 1-step and 2-step correction. An ingress port is supposed to accept both 1-step or 2-step corrected messages; it does not need special hardware for this.

A TC forwards the message with 1-step or 2-step correction, possibly modifying the 2-step flag in the message header. A BC treats the clock domains independently.

Although IEC 61588:2009 lists the ability of performing 1-step or 2-step as a default clock property, it is in reality a port property, but usually the same hardware will be used on different ports, so IEC 61588:2009 assumes the same property for all ports.

D.6.2.2 Translation from 1-step to 2-step and vice-versa

A 2-step clock is always able to receive 1-step Sync messages. A 2-step clock that receives a 1-step Sync needs no special translation except change the mode to 2-step and generate the Follow_Up. In one method, the Sync carries the received correction field, while the Follow_Up carries the sum of the residence delay and of the peer delay. Alternatively, the Sync carries the sum of the received correction field and of the peer delay, while the Follow_Up carries only the residence delay. The next clock will take the sum of the correction fields in the Sync and in the Follow_Up messages for correction, as IEC 61588:2009 foresees.

NOTE A 1-step TC modifies the correction field of the Sync message, while a 2-step TC modifies the correction field of the Follow_Up message. Therefore it is important that for any 1-step translation or synchronization the sum of the Sync and Follow_Up correction field is used.

D.6.2.3 Translation from 2-step to 1-step

NOTE 1 IEC 61588:2009 only foresees a 1-step to 2-step translation. The method described in D.6.2.3 is considered for a next revision of IEC 61588, it is fully backward compatible.

Figure D.6 shows the translation from 2-step to 1-step in a chain of 1-step and 2-step TCs between an ordinary clock (slave) and the grandmaster clock (master).





Figure D.6 – Translation from 2-step to 1-step in TCs

A 1-step clock that receives a 2-step Sync awaits the Follow_Up before generating the 1-step Sync. The 1-step Sync at the egress port then carries the sum of the correction fields of the 2-step Sync, of the Follow_Up, the residence delay of the TC and the estimated delay to the upstream peer, as Figure D.7 shows.

NOTE 2 Without synchronized clocks the error can be large. A 500 ms delay between Sync and Follow_Up measured by a 50 ppm local oscillator would cause an error of 25 μ s.



IEC

Figure D.7 – Translation from 2-step to 1-step – message view

D.7 End-To-End link delay measurement

D.7.1 General method

In all methods, the residence time of a Sync message in the TCs is measured using the local oscillator, the difference is in the method for measuring the link delays.

In end-to-end link delay measurement, the TCs correct the Sync or Follow_Up messages only for the residence time, but not for the link delays. The calculation of the link delay involves a call-and-reply mechanism using the delay request (Delay_Req) and delay response (Delay_Resp) messages, which yields the sum of all link delays along the path that Sync takes, assuming that link delays are the same in both directions.

To minimize the time error due to the clock drift, the time difference between Sync and Delay_Req is kept short. In practice, Delay_Req is sent in response to a Sync and is evaluated for the next Sync. Path delays do not change rapidly, so this value can be averaged over several periods.

The handling of these messages is different in 1-step and 2-step correction.

D.7.2 End-to-End link delay measurement with 1-step clock correction

Each slave calculates the sum of the link delays on the path to the master by sending a Delay_Req (delay request) message to the master, which responds with a Delay_Resp (delay response) message containing the time of reception. Both Delay_Req and Delay_Resp cross all TCs in the network. The slave awaits the Sync to compute the correction.

The master sends the Sync at time t_1 and the TCs correct the received value of t_1 by their residence time ρ . The slave receive the Sync at time t_2 with the sum of the residence time correction ρ_{ms} .

To calculate the sum of the end-to-end link delays, the slave sends at time t_3 a Delay_Req to the master. The master receives the Delay_Req at time t_4 with the sum of the residence time corrections on the path and responds with a Delay_Resp that contains the received sum of the correction and the timestamp t_4 , as Figure D.8 shows.




Figure D.8 – End-to-end link delay measurement with 1-step clock correction

After this exchange, the slave keeps t_3 , t_4 and the sum of the residence times, $\Sigma \rho_{sm}$.

TCs with 1-step correct the timestamp in the Delay_Req for their residence delay. They do not modify the Delay_Resp.

This mode therefore implies that the Delay_Req and Sync take the same path, which requires some engineering in a layer 3 network. The path taken by Delay_Resp is less critical.

D.7.3 End-to-End link delay measurement with 2-step clock correction

TCs with 2-step correction do not modify the Delay_Req; they just record the exit timestamp of the Delay_Req. When the Delay_Resp comes back, they insert this correction when they forward the Delay_Resp in the opposite direction, as Figure D.9 shows. The Delay_Resp collects all residence delays.

This mode therefore implies that the TCs memorize the communicating partners and that Sync, Delay_Req and Delay_Resp take the same path, which requires some engineering in a layer 3 network.




Figure D.9 – End-to-end delay measurement with 2-step clock correction

D.7.4 End-to-End link delay calculation by Delay_Req/Delay_Res

In both 1-step and 2-step correction, the slave deduces the sum of the two link delays from the correction field $\Sigma\delta_{sm}$ and $\Sigma\rho_{ms}$ previously received in Delay_Res and Sync:

$$\Sigma\lambda = \frac{(t_1 - t_4) - (t_2 - t_3) - \Sigma\rho_{ms} - \Sigma\rho_{sm}}{2}$$

Again, this scheme assumes symmetry of the path delays (the residence times can be widely different).

D.8 Peer-to-Peer link delay calculation


D.8.1 Peer-to-Peer link delay calculation with 1-step correction

Unlike end-to-end which measures the sum of the link delays, each peer-to-peer TC computes the link delay to all its neighbours by sending a time-stamped Pdelay_Req (peer delay request) message on its links, to which the peer responds with a time-stamped Pdelay_Res (peer delay response) message. This delay is summed with the residence delay in the correction of the Sync or Follow_Up.

Figure D.10 shows a PTP clock distribution peer-to-peer link delay measurement and with 1-step correction.

To calculate the link delays, each clock sends peer delay request messages (Pdelay_Req) to all its neighbours, even in the direction opposite to the current GMC since a redundant GMC may exist. The neighbour responds with a Pdelay_Resp message and the link delay over each link is calculated as:

$$\lambda = \frac{(t_4 - t_1) - (t_3 - t_2)}{2}$$



IEC

Figure D.10 – Peer-to-peer link delay measurement with 1-step clock correction

In 1-step clocks, TC_i adds the received correction from clock *k* to its own correction ($\lambda_{kj} + \rho_j$) and modifies the Sync message's correction field on-the-fly, while the frame is being transmitted over the egress port, since the residence time is only precisely known when transmission started.

This 1-step correction requires a dedicated hardware to modify the correction field (and the checksum) of the outgoing Sync frame while the header of the message has already been transmitted.

The slave clock corrects the received time t_4 by the value of the correction field, which contains the sum of the residence times and the sum of the estimated link delays.

D.8.2 Peer-to-Peer link delay calculation with 2-step correction

To avoid a dedicated hardware for modifying the Sync message on-the-fly, a 2-steps clock records the precise time at which the Sync was received, forwards the Sync message without modification (but announcing a Follow_Up message) and records the precise time at which the Sync was sent. Dedicated PHY chips time-stamp messages within a few nanoseconds, so sending the Sync message is not time-critical, it can be done by a CPU.

The 2-step TCs calculates the correction field as the sum of the ingress Sync correction, the residence delay and the (previously measured) peer delay and sends a Follow_Up message as Figure D.11 shows. The same scheme applies to the calculation of the peer delay.

The receiver assumes that Sync and Follow_Up messages are paired; they have the same PTP sequence number.




Figure D.11 – Peer-to-peer link delay measurement with 2-step clock correction

A 2-step Pdelay_Resp/Pdelay_Resp_Follow_Up improves over 1-step Pdelay_Resp since it transmits the values of t_2 and t_3 individually and therefore allows to discover errors in the syntonization, as described in IEC 61588:2009, 6.5.5 and specified in IEC 61588:2009, 11.4.3, (c) 8) option (i).

A combination of 2-step with 1-step is in principle possible, since the 2-step flag applies to the messages individually, but is not currently specified. In view of a possible revision of IEC 61588:2009, the receiver should watch the 2-step flag of the received Pdelay_Resp or Sync messages to know if a Pdelay_Resp_Follow_Up or Follow_Up is expected. Also, the revision of IEC 61588:2009 is expected to define a time-out for a Follow_Up to be received, in case of unreliable links.

Annex E (normative)

Management Information base for singly and doubly attached clocks

NOTE When copying the MIB from a PDF document, all headers and footers are removed and characters are converted to the character format of SNMP. Although care has been taken not to use illegal characters, some can appear in the final document due to the editing process.

```

IEC62439-8-MIB DEFINITIONS ::= BEGIN
IMPORTS
    MODULE-IDENTITY, OBJECT-TYPE, NOTIFICATION-TYPE, Unsigned32, Integer32, Counter32
FROM SNMPv2-SMI -- [RFC2578]
    TEXTUAL-CONVENTION, TruthValue FROM SNMPv2-TC -- [RFC2579]
    MODULE-COMPLIANCE, OBJECT-GROUP, NOTIFICATION-GROUP FROM SNMPv2-CONF --
[RFC2580]
    iec62439 FROM IEC-62439-3-MIB -- root
;
ptp MODULE-IDENTITY
    LAST-UPDATED "201504081548Z"-- 2015-04-08: 15:48:55
    ORGANIZATION "IEC/SC 65C WG15"
    CONTACT-INFO "International Electrotechnical Commission,
Central Office, 3, rue de Varambe 22, CH-1211 Geneva 20, Switzerland"
    DESCRIPTION
        "The Management Information Base module of the IEC62439 Redundant Time
Distribution protocol."
    REVISION "201504081548Z"-- 2015-04-08: 15:48:55
    DESCRIPTION
        "Original version to be published as part of the CDV"
        ::= { iso(1) std(0) 62439 }
    OBJECT IDENTIFIER ::= { iec62439.ptp(7) }

=====
-- Textual Conventions
=====

ClockIdentity ::= TEXTUAL-CONVENTION
    DISPLAY-HINT "1x:"
    STATUS current
    DESCRIPTION
        "Represents an IEEE 802 MAC address represented in the
canonical order defined by IEEE 802.1a, EUI-64. EUI-48
converts to EUI-64 as specified by IEEE. The conversion
assigns values 255 and 254 to octets 3 and 4 respectively,
where octet 0 is the most significant and octet 7 the least.
For example, EUI-48 of AC:DE:48:23:45:67 would extend to
AC:DE:48:FF:FE:23:45:67. In IEC 62439-3 the octets 4 and 5 are used to transport the
GMshortID"
    REFERENCE "IEC 61588:2009 5.3.4 and 7.5.2.2"
    SYNTAX OCTET STRING (SIZE (8))

PtpTimeInterval ::= TEXTUAL-CONVENTION
    DISPLAY-HINT "2x:"
    STATUS current
    DESCRIPTION
        "TimeInterval Value from IEEE Std 1588-2008 5.3.2 - 2**48 / 10**9 s"
    REFERENCE "IEC 61588:2009 5.3.2"
    SYNTAX OCTET STRING (SIZE (8))

ClockClass ::= TEXTUAL-CONVENTION
    STATUS current
    DESCRIPTION
        "Clock Class Value from IEEE Std 1588-2008 7.6.2.4,
with the following interpretation placed on the value:

```

6: A clock that is synchronized to a primary reference time source,
 7: A clock that has previously been designated as clockClass 6 but that has lost the ability to synchronize to a primary reference time source and is in holdover mode and within holdover specifications,
 13: A clock that is synchronized to an application-specific source of time,
 14: A clock that has previously been designated as clockClass 13 but that has lost the ability to synchronize to an application-specific source of time and is in holdover mode and within holdover specifications,
 52: Degradation alternative A for a clock of clockClass 7 that is not within holdover specification,
 58: Degradation alternative A for a clock of clockClass 14 that is not within holdover specification,
 68..122: For use by alternate PTP profiles (68..122),
 133..170: For use by alternate PTP profiles (133..170),
 187: Degradation alternative B for a clock of clockClass 7 that is not within holdover specification,
 193: Degradation alternative B for a clock of clockClass 14 that is not within holdover specification,
 216..232: For use by alternate PTP profiles,
 248: Default none of the other clockClass definitions apply,
 255: A slave-only clock(255)."

REFERENCE "IEC 61588:2009 7.6.2.4"

SYNTAX INTEGER {
 primarySync(6),
 primarySyncLost(7),
 applicationSpecificSync(13),
 applicationSpecificSyncLost(14),
 primarySyncAlternativeA(52),
 applicationSpecificAlternativeA(58),
 primarySyncAlternativeB(187),
 applicationSpecificAlternativeB(193),
 default(248),
 slaveOnlyClock(255)
 }

ClockAccuracy ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Clock Accuracy Value from IEEE Std 1588-2008 8.6.2.3, with the following

interpretation placed on the value:

32: The time is accurate to within 25 ns,
 33: The time is accurate to within 100 ns,
 34: The time is accurate to within 250 ns,
 35: The time is accurate to within 1 us,
 36: The time is accurate to within 2.5 us,
 37: The time is accurate to within 10 us,
 38: The time is accurate to within 25 us,
 39: The time is accurate to within 100 us,
 40: The time is accurate to within 250 us,
 41: The time is accurate to within 1 ms,
 42: The time is accurate to within 2.5 ms,
 43: The time is accurate to within 10 ms,
 44: The time is accurate to within 25 ms,
 45: The time is accurate to within 100 ms,
 46: The time is accurate to within 250 ms,
 47: The time is accurate to within 1 s,
 48: The time is accurate to within 10 s,
 49: The time is accurate to within > 10 s,
 254: Default indicating unknown"

REFERENCE "IEC 61588:2009 7.6.2.5 and Table 6"

SYNTAX INTEGER {
 timeAccurateTo25ns(32),


```

timeAccurateTo100ns(33),
timeAccurateTo250ns(34),
timeAccurateTo1000ns(35),
timeAccurateTo2500ns(36),
timeAccurateTo10000ns(37),
timeAccurateTo25000ns(38),
timeAccurateTo100000ns(39),
timeAccurateTo250000ns(40),
timeAccurateTo1000000ns(41),
timeAccurateTo2500000ns(42),
timeAccurateTo10ms(43),
timeAccurateTo25ms(44),
timeAccurateTo100ms(45),
timeAccurateTo250ms(46),
timeAccurateTo1s(47),
timeAccurateTo10s(48),
timeAccurateToGT10s(49),
timeAccurateToUnknown(254)
}

```

TimeSourceValue ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"The timeSource is an information only attribute indicating the type of source of time used by a grandmaster clock.

The following value are defined:

```

16: Atomic Clock,
32: GPS,
48: Terrestrial Radio,
64: PTP,
80: NTP,
96: Hand Set,
112: IRIG-B,
144: Other,
160: Internal Oscillator"

```

REFERENCE "IEC 61588:2009 7.6.2.6 and Table 7"

```

SYNTAX INTEGER {
    atomicClock(16),
    gps(32),
    terrestrialRadio(48),
    ptp(64),
    ntp(80),
    handSet(96),
    irigb(112),
    other(144),
    internalOscillator(160)
}

```

Integer48 ::= TEXTUAL-CONVENTION

DISPLAY-HINT "12d"

STATUS current

DESCRIPTION

"48bit Integer value"

REFERENCE "IEEE 802.3"

SYNTAX OCTET STRING (SIZE (6))

DelayMechanism ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Identifies the path delay measurement mechanism in use"

REFERENCE "IEC 61588:2009 8.2.5.4.4"

```

SYNTAX INTEGER {
    e2e(1),
    p2p(2),
    disabled(254)
}

```

PortStateType ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"The state of the port protocol machine according to 1588 "

REFERENCE "IEC 61588:2009 8.2.5.4.4"

SYNTAX INTEGER {
 undefined (0),
 initializing (1),
 faulty (2),
 disabled (3),
 listening (4),
 preMaster (5),
 master (6),
 passiveMaster (7),
 uncalibrated (8),
 slave (9),
 passiveSlave (10)
 }

ProfileSet ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"Identifies the profile supported or in use"

REFERENCE "IEC 62439-3"

SYNTAX BITS {
 resProfile0 (0),
 powerUtilityProfile (7),
 l3e2eProfile (8),
 l2p2pProfile (9)
 }

PortAttachmentType ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"port attachment type according to IEC 62439-3"

REFERENCE "IEC 62439-3"

SYNTAX INTEGER {
 unspecified (0),
 ordinaryClockPort (1),
 boundaryClockPort (2),
 transparentClockPort (3),
 boundaryNode (4),
 doublyAttachedOcPort (5),
 doublyAttachedBcPort (6),
 doublyAttachedTcPort (7),
 pairedStatelessTcPort (8),
 other (255)
 }

TlvTxRxIndex ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"index for transmitted or received TLV object"

REFERENCE "IEC 62439-3"

SYNTAX INTEGER {
 tx (1),
 rx (2)
 }

NetworkProtocol ::= TEXTUAL-CONVENTION

STATUS current

DESCRIPTION

"PTP transport protocol in use over that port"

REFERENCE "IEC 61588:2009 7.4.1 Table 3"

SYNTAX INTEGER {

```

    reserved (0),
    udpIpv4 (1),
    udpIPv6 (2),
    ieee8023 (3),
    deviceNet (4),
    controlNet (5),
    profinet (6),
    reserved2 (65535)
}

```

LogTimeInterval ::= TEXTUAL-CONVENTION

DISPLAY-HINT "d"

STATUS current

DESCRIPTION

"Time interval expressed as logarithm base 2 of the number of seconds, e.g. 0 = 1s, 2 = 4s, -3 = 125ms, "

REFERENCE "IEC 61588:2009 7.7.2.1"

SYNTAX Integer32 (-128..127)

```

=====
-- Events
=====

```

ptpNotifications OBJECT IDENTIFIER ::= { ptp 0}

eventFaultyState NOTIFICATION-TYPE

OBJECTS { ptpPortDsPortState }

STATUS current

DESCRIPTION

"A clock detects it has entered a faulty state."

::= { ptpNotifications 1 }

eventRedundancyLoss NOTIFICATION-TYPE

OBJECTS { ptpPortDsPortState }

STATUS current

DESCRIPTION

"A clock that was previously redundantly connected detects loss of one channel"

::= { ptpNotifications 2 }

eventPeerDelayExceeded NOTIFICATION-TYPE

OBJECTS { ptpPortDsPeerDelayLim }

STATUS current

DESCRIPTION

"A port detects an abnormal increase of the peer delay, possibly due to link failure"

::= { ptpNotifications 3 }

```

=====
-- Data Sets
=====

```

ptpObjects OBJECT IDENTIFIER ::= { ptp 1}

ptpOcBc OBJECT IDENTIFIER ::= { ptpObjects 1}

ptpDefaultDs OBJECT IDENTIFIER ::= { ptpOcBc 1}

ptpDefaultDsTwoStepFlag OBJECT-TYPE

SYNTAX TruthValue

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"If all ports of the clock apply two-step time correction: true; otherwise false (The port settings may allow individual two-step correction)"

REFERENCE "IEC61588 8.2.1.2.1"

DEFVAL {true}

::= { ptpDefaultDs 1 }

```
-- .1.1.1

ptpDefaultDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The grandMasterIdentity in the case this clock becomes grandmaster
The octets 4 and 5 should carry the value 0xFFFE."
    REFERENCE   "IEC61588 8.2.1.2.2, 7.6.2.1, 7.5.2.2"
    ::= { ptpDefaultDs 2 }
-- .1.1.2

ptpDefaultDsNumberPorts OBJECT-TYPE
    SYNTAX      Unsigned32(1..255)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The number of PTP ports on the device. (1 if OC, >1 for Boundary Clock).
IEC 62439-3 allows 2 for OC that are doubly attached"
    REFERENCE   "IEC61588 8.2.1.2.3, IEC62439-3 A.8.3.2.1"
    ::= { ptpDefaultDs 3 }
-- .1.1.3

ptpDefaultDsClockClass OBJECT-TYPE
    SYNTAX      ClockClass
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The clockClass specification"
    REFERENCE   "IEC61588 8.2.1.3.1.1, 7.6.2.4."
    ::= { ptpDefaultDs 4 }
-- .1.1.4

ptpDefaultDsClockAccuracy OBJECT-TYPE
    SYNTAX      ClockAccuracy
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The clockAccuracy member of the ptpDefaultDs.clockQuality member"
    REFERENCE   "IEC61588 8.2.1.3.1.2, 5.3.7"
    ::= { ptpDefaultDs 5 }
-- .1.1.5

ptpDefaultDsOffsetScaledLogVaria OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The offsetScaledLogVariance member of the ptpDefaultDs.clockQuality member"
    REFERENCE   "IEC61588 8.2.1.3.1.3, 5.3.7 and 7.6.3.5"
    ::= { ptpDefaultDs 6 }
-- .1.1.6

ptpDefaultDsPriority1 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The priority1 attribute of the local clock"
    REFERENCE   "IEC61588 8.2.1.4.1, 7.6.2.2"
    ::= { ptpDefaultDs 7 }
-- .1.1.7

ptpDefaultDsPriority2 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
```

```

        "The priority2 attribute of the local clock"
    REFERENCE      "IEC61588 8.2.1.4.2, 7.6.2.3"
    ::= { ptpDefaultDs 8 }
-- .1.1.8

ptpDefaultDsDomainNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The domain attribute of the local clock, it applies to all ports in a BC"
    REFERENCE      "IEC61588 8.2.1.4.3, 7.1"
    ::= { ptpDefaultDs 9 }
-- .1.1.9

ptpDefaultDsSlaveOnly OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "If the clock is a slave-only clock: true; otherwise false"
    REFERENCE      "IEC61588 8.2.1.4.4, 9.2.2. and 9.2.3"
    ::= { ptpDefaultDs 10 }
-- .1.1.10

ptpDefaultDsProfileSet OBJECT-TYPE
    SYNTAX      ProfileSet
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The set of profiles which the clock has the ability to support with the
        required performance. Profile activation is a port property"
    REFERENCE      "IEC62439-3 A.8.3.2.2"
    ::= { ptpDefaultDs 11 }
-- .1.1.11

ptpDefaultDsTimeInaccuracy OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The absolute value of the rated worst-case time error with respect to
        absolute time when the clock is connected to its reference signal. If the clock is a
        boundary clock, the reference signal is the clock signal applied to any port, and the
        rated time error the error observed on any master port."
    REFERENCE      "IEC62439-3 A.8.3.2.3"
    ::= { ptpDefaultDs 12 }
-- .1.1.12

ptpDefaultDsOffsetFromMasterLim OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The value of the limit to the offset from master that signals an abnormal
        condition"
    REFERENCE      "IEC62439-3 A.8.3.2.4"
    ::= { ptpDefaultDs 13 }
-- .1.1.13

ptpCurrentDs OBJECT IDENTIFIER ::= { ptpOcBc 2}

ptpCurrentDsStepsRemoved OBJECT-TYPE
    SYNTAX      Unsigned32
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION

```

```

        "The number of boundary clocks traversed between the local clock and the
grandmaster clock"
        REFERENCE      "IEC61588 8.2.2.2"
        ::= { ptpCurrentDs 1 }
-- .1.2.1

ptpCurrentDsOffsetFromMaster OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS       "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The implementation-specific, current value of the time difference between a
master and a slave as computed by the slave"
        REFERENCE      "IEC61588 8.2.2.3"
        ::= { ptpCurrentDs 2 }
-- .1.2.2

ptpCurrentDsMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS       "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The implementation-specific, current value of the mean propagation time
between a master and slave clock as computed by the slave"
        REFERENCE      "IEC61588 8.2.2.4"
        ::= { ptpCurrentDs 3 }
-- .1.2.3

ptpParentDs OBJECT IDENTIFIER ::= { ptpOcBc 3}

ptpParentDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The identity of the port on the master that issues the Sync messages"
        REFERENCE      "IEC61588 8.2.3.2"
        ::= { ptpParentDs 1 }
-- .1.3.1

ptpParentDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The number of the port on the master that issues the Sync"
        REFERENCE      "IEC61588 8.2.3.2"
        ::= { ptpParentDs 2 }
-- .1.3.2

ptpParentDsStats OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "If all of the following conditions are satisfied:
- The clock has a port in the SLAVE state.
- The clock has computed statistically valid estimates of the
ptpParentDs.observedParentOffsetScaledLog Variance and the
ptpParentDs.observedParentClockPhaseChangeRate members: true
otherwise false."
        REFERENCE      "IEC61588 8.2.3.3"
        ::= { ptpParentDs 3 }
-- .1.3.3

ptpParentDsObsOfstScdLVar OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only

```

```

    STATUS      current
    DESCRIPTION
        "If ptpParentDs.parentStats is true, an estimate PTP variance of the parent
        clock as observed by the slave clock. The computation of this value is optional;
        If ptpParentDs.parentStats is false: not valid"
    REFERENCE   "IEC61588 8.2.3.4, 7.6.3.5, 7.6.3.3"
    ::= { ptpParentDs 4 }
-- .1.3.4

ptpParentDsObsPhChgRate OBJECT-TYPE
    SYNTAX      Integer32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The estimate of the phase change rate of the parent clock
        (ptpParentDs.observedParentClockPhaseChangeRate) as observed by the slave clock. If
        the estimate exceeds the capacity of its data type, this value is set to 7FFF FFFF16
        or 8000 000016, as appropriate."
    REFERENCE   "IEC61588 8.2.3.5, 7.6.4.4, and 15.5.3.5.1.4"
    ::= { ptpParentDs 5 }
-- .1.3.5

ptpParentDsGmIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The received grandMasterIdentity"
    REFERENCE   "IEC61588 8.2.3.6, 7.5.2.2, and 5.3.4"
    ::= { ptpParentDs 6 }
-- .1.3.6

ptpParentDsGmClass OBJECT-TYPE
    SYNTAX      ClockClass
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The clockQuality.clockClass attribute of the grandmaster clock."
    REFERENCE   "IEC61588 8.2.3.7, 7.6.2.4"
    ::= { ptpParentDs 7 }
-- .1.3.8

ptpParentDsGmAccuracy OBJECT-TYPE
    SYNTAX      ClockAccuracy
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The clockQuality.clockAccuracy attribute of the grandmaster clock."
    REFERENCE   "IEC61588 8.2.3.7, 7.6.2.5"
    ::= { ptpParentDs 8 }
-- .1.3.9

ptpParentDsGmOfstScdLVar OBJECT-TYPE
    SYNTAX      Integer32
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The clockQuality.offsetScaledLogVariance attribute of the grandmaster
        clock."
    REFERENCE   "IEC61588 8.2.3.7, 7.6.3, 7.6.2.2"
    ::= { ptpParentDs 9 }
-- .1.3.10

ptpParentDsGmPriority1 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS  read-only
    STATUS      current
    DESCRIPTION
        "The priority1 attribute of the grandmaster clock."
    REFERENCE   "IEC61588 8.2.3.8, 7.6.2.2"

```

```

    ::= { ptpParentDs 10 }
-- .1.3.11

ptpParentDsGmPriority2 OBJECT-TYPE
    SYNTAX      Unsigned32(0..255)
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The priority2 attribute of the grandmaster clock."
    REFERENCE    "IEC61588 8.2.3.9, 7.6.2.3"
    ::= { ptpParentDs 11 }
-- .1.3.12

ptpTimePropDs OBJECT IDENTIFIER ::= { ptpOcBc 4}

ptpTimePropDsCurrentUtcOffset OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and if the epoch is the PTP epoch, the
value of defaultDS.currentUtcOffset (offset between TAI and UTC in seconds);
otherwise the value has no meaning"
    REFERENCE    "IEC61588 8.2.4.2, 7.2.3, 9.4"
    ::= { ptpTimePropDs 1 }
-- .1.4.1

ptpTimePropDsCurrentUtcOffsetVld OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and if the defaultDS.currentUtcOffset is
known to be correct: true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.3, 7.2.3, 9.4"
    ::= { ptpTimePropDs 2 }
-- .1.4.2

ptpTimePropDsLeap59 OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and the epoch is the PTP epoch and the last
minute of the current UTC day contains 59 seconds: true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.4, 9.4"
    ::= { ptpTimePropDs 3 }
-- .1.4.3

ptpTimePropDsLeap61 OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the clock class is 6, 7 or 52 and the epoch is the PTP epoch and the last
minute of the current UTC day contains 61 seconds (timePropertiesDS.leap61): true;
otherwise false"
    REFERENCE    "IEC61588 8.2.4.5, 9.4"
    ::= { ptpTimePropDs 4 }
-- .1.4.4

ptpTimePropDsTimeTraceable OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "If the timescale is PTP and the source is traceable to a primary reference:
true; otherwise false"
    REFERENCE    "IEC61588 8.2.4.6"
    ::= { ptpTimePropDs 5 }

```


-- .1.4.5

ptpTimePropDsFrequencyTraceable OBJECT-TYPE

SYNTAX TruthValue

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"If the frequency determining the timescale is traceable to a primary reference: true; otherwise false"

REFERENCE "IEC61588 8.2.4.7"

::= { ptpTimePropDs 6 }

-- .1.4.6

ptpTimePropDsPtpTimescale OBJECT-TYPE

SYNTAX TruthValue

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"If the clock timescale of the grandmaster clock is PTP: true; otherwise false"

REFERENCE "IEC61588 8.2.4.8, 7.2.1"

::= { ptpTimePropDs 7 }

-- .1.4.7

ptpTimePropDsTimeSource OBJECT-TYPE

SYNTAX TimeSourceValue

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"The source of time used by the grandmaster clock given by the enumeration type"

REFERENCE "IEC61588 8.2.4.9"

::= { ptpTimePropDs 8 }

-- .1.4.8

ptpForeignMasterDsTable OBJECT-TYPE

SYNTAX SEQUENCE OF PtpForeignMasterDsEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A table of foreign master entries"

REFERENCE "IEC61588 9.3.2"

::= { ptpOcBc 5 }

-- .1.5

ptpForeignMasterDsEntry OBJECT-TYPE

SYNTAX PtpForeignMasterDsEntry

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"A list of foreign master entries"

REFERENCE "IEC61588 9.3.2"

INDEX {
ptpForeignMasterDsIndex }

::= { ptpForeignMasterDsTable 1 }

-- .1.5.1

PtpForeignMasterDsEntry ::= SEQUENCE {

ptpForeignMasterDsIndex Unsigned32,

ptpForeignMasterDsPortIdentity ClockIdentity,

ptpForeignMasterAnnounceMessages Unsigned32

}

ptpForeignMasterDsIndex OBJECT-TYPE

SYNTAX Unsigned32(0..65535)

MAX-ACCESS not-accessible

STATUS current

DESCRIPTION

"Index for ptpForeignMasterDs Table"

```

REFERENCE      "IEC61588 9.3.2.4.5"
::= { ptpForeignMasterDsEntry 1 }
-- .1.5.1.1

ptpForeignMasterDsPortIdentity OBJECT-TYPE
SYNTAX      ClockIdentity
MAX-ACCESS   read-only
STATUS       current
DESCRIPTION
    "The sourcePortIdentity field value of an Announce message received from the
    foreign master"
REFERENCE      "IEC61588 9.3.2.4.2"
::= { ptpForeignMasterDsEntry 2 }
-- .1.5.1.2

ptpForeignMasterAnnounceMessages OBJECT-TYPE
SYNTAX      Unsigned32
MAX-ACCESS   read-only
STATUS       current
DESCRIPTION
    "The number of Announce messages from the foreign master indicated by the
    foreignMasterDS.foreignMasterPortIdentity member
    that have been received within a time window of duration FOREIGN_MASTER_TIME_WINDOW"
REFERENCE      "IEC61588 9.3.2.4.3"
::= { ptpForeignMasterDsEntry 3 }
-- .1.5.1.2

ptpPortDsTable OBJECT-TYPE
SYNTAX      SEQUENCE OF PtpPortDsEntry
MAX-ACCESS   not-accessible
STATUS       current
DESCRIPTION
    "A table of port related variables, a value of 1 is used in devices with only
    one port"
::= { ptpOcBc 6 }
-- .1.6

ptpPortDsEntry OBJECT-TYPE
SYNTAX      PtpPortDsEntry
MAX-ACCESS   not-accessible
STATUS       current
DESCRIPTION
    "A list of objects pertaining to a port."
REFERENCE      "IEC61588 8.2.5"
INDEX       {
    ptpPortDsIndex }
::= { ptpPortDsTable 1 }
-- .1.6.1

PtpPortDsEntry ::= SEQUENCE {
    ptpPortDsIndex      Unsigned32,
    ptpPortDsClockIdentity ClockIdentity,
    ptpPortDsPortNumber Unsigned32,
    ptpPortDsPortState   PortStateType,
    ptpPortDsLogMinDelayReqInterval LogTimeInterval,
    ptpPortDsPeerMeanLinkDelay PtpTimeInterval,
    ptpPortDsLogAnnounceInterval LogTimeInterval,
    ptpPortDsAnnounceReceiptTimeout Unsigned32,
    ptpPortDsLogSyncInterval LogTimeInterval,
    ptpPortDsDelayMechanism DelayMechanism,
    ptpPortDsLogMinPdelayReqInterval LogTimeInterval,
    ptpPortDsVersionNumber Unsigned32,
    ptpPortDsNetworkProtocol NetworkProtocol,
    ptpPortDsPortEnabled TruthValue,
    ptpPortDsDlyAsymmetry PtpTimeInterval,
    ptpPortDsProfileId ProfileSet,
    ptpPortDsVlanEnable TruthValue,
    ptpPortDsVlanId Unsigned32,
    ptpPortDsVlanPrio Unsigned32,

```

```

        ptpPortDsTwoStepFlag      TruthValue,
        ptpPortDsPeerIdentity     ClockIdentity,
        ptpPortDsPrpAttachement   PortAttachmentType,
        ptpPortDsPrpPairedPort    Unsigned32,
        ptpPortDsErrorCounter     Counter32,
        ptpPortDsPeerDelayLim     PtpTimeInterval
    }

ptpPortDsIndex OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   not-accessible
    STATUS      current
    DESCRIPTION
        "Index for ptpPortDs Table"
    ::= { ptpPortDsEntry 1 }
-- .1.6.1.1

ptpPortDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The portIdentity.clockIdentity attribute of the local port"
    REFERENCE   "IEC61588 8.2.5.2.1, 7.5.2.2.1"
    ::= { ptpPortDsEntry 2 }
-- .1.6.1.2

ptpPortDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The portIdentity.portNumber attribute of the local port"
    REFERENCE   "IEC61588 8.2.5.2.1, 7.5.2.3"
    ::= { ptpPortDsEntry 3 }
-- .1.6.1.3

ptpPortDsPortState OBJECT-TYPE
    SYNTAX      PortStateType
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The current state of the protocol engine associated with this port taken
        from the enumeration in Table 8, extended by the PASSIVE_SLAVE state in doubly
        attached clocks"
    REFERENCE   "IEC61588 8.2.5.3.1, 9.2; IEC62439-3 A.8.3.4.1"
    ::= { ptpPortDsEntry 4 }
-- .1.6.1.4

ptpPortDsLogMinDelayReqInterval OBJECT-TYPE
    SYNTAX      LogTimeInterval
    UNITS        "seconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The logarithm to the base 2 of the minDelayReqInterval; The initialization
        value is implementation-specific consistent with 7.7.2.4."
    REFERENCE   "IEC61588 8.2.5.3.2, 7.7.2.4"
    ::= { ptpPortDsEntry 5 }
-- .1.6.1.5

ptpPortDsPeerMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION

```

"If the value of the ptpPortDs.delayMechanism member is peer-to-peer (P2P), an estimate of the current one-way propagation delay on the link, i.e., <meanPathDelay>, attached to this port computed using the peer delay mechanism. If the value of the ptpPortDs.delayMechanism member is end-to-end (E2E), this value is zero"

```
REFERENCE      "IEC61588 8.2.5.3.3, 11.4"
::= { ptpPortDsEntry 6 }
-- .1.6.1.6
```

ptpPortDsLogAnnounceInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the mean announceInterval"
REFERENCE   "IEC61588 8.2.5.4.1, 7.7.2.2"
::= { ptpPortDsEntry 7 }
-- .1.6.1.7
```

ptpPortDsAnnounceReceiptTimeout OBJECT-TYPE

```
SYNTAX      Unsigned32(0..255)
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The integral multiple of announceInterval"
REFERENCE   "IEC61588 8.2.5.4.2, 7.7.3.1"
::= { ptpPortDsEntry 8 }
-- .1.6.1.8
```

ptpPortDsLogSyncInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the mean SyncInterval for multicast messages"
REFERENCE   "IEC61588 8.2.5.4.3, 7.7.2.3"
::= { ptpPortDsEntry 9 }
-- .1.6.1.9
```

ptpPortDsDelayMechanism OBJECT-TYPE

```
SYNTAX      DelayMechanism
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The propagation delay measuring option used by the port in computing
    <meanLinkDelay>, the data type being enumeration taken from Table 9. The
    initialization value is implementation-specific unless otherwise stated."
REFERENCE   "IEC61588 8.2.5.4.4"
::= { ptpPortDsEntry 10 }
-- .1.6.1.10
```

ptpPortDsLogMinPdelayReqInterval OBJECT-TYPE

```
SYNTAX      LogTimeInterval
UNITS       "seconds"
MAX-ACCESS  read-write
STATUS      current
DESCRIPTION
    "The logarithm to the base 2 of the minPdelayReqInterval"
REFERENCE   "IEC61588 8.2.5.4.5, 7.7.2.5"
::= { ptpPortDsEntry 11 }
-- .1.6.1.11
```

ptpPortDsVersionNumber OBJECT-TYPE

```
SYNTAX      Unsigned32(2..15)
MAX-ACCESS  read-only
STATUS      current
DESCRIPTION
```

```

        "The PTP version in use on the port; it is 2 for IEC 61588-2009"
    REFERENCE      "IEC61588 8.2.5.4.6"
    ::= { ptpPortDsEntry 12 }
-- .1.6.1.12

ptpPortDsNetworkProtocol OBJECT-TYPE
    SYNTAX      NetworkProtocol
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The network protocol in use over that port according to Table 3"
    REFERENCE      "IEC61588 7.4.1"
    ::= { ptpPortDsEntry 13 }
-- .1.6.1.13

ptpPortDsPortEnabled OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "Enables the port for PTP traffic and indicates the enabled status of the
port if true"
    REFERENCE      "IEC62439-3 A.8.3.4.2"
    DEFVAL       {1}
    ::= { ptpPortDsEntry 14 }
-- .1.6.1.14

ptpPortDsDlyAsymmetry OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The link delay asymmetry to correct the link delay"
    REFERENCE      "IEC61588 7.4.2, IEC62439-3 A.8.3.4.3"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpPortDsEntry 15 }
-- .1.6.1.15

ptpPortDsProfileId OBJECT-TYPE
    SYNTAX      ProfileSet
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The profiles in use on this port"
    REFERENCE      "IEC62439-3 A.8.3.4.4"
    ::= { ptpPortDsEntry 16 }
-- .1.6.1.16

ptpPortDsVlanEnable OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "Enables that PTP messages generated internally to the clock receive an
802.1Q tag with the values specified in portDs.vlanId and portDs.vlanPrio when sent
over this port, with exception of PdelayReq and PdelayResp, which operate with link-
local addresses. It does not apply to messages in transit generated by other clocks.
The bridge rules take precedence over this setting.
This does not concern PTP messages in transit (even modified), which are handled
according to the bridging rules.
In a boundary clock, this setting applies to the sending over that port only."
    REFERENCE      "IEC62439-3 A.8.3.4.5; IEEE 802.1Q"
    ::= { ptpPortDsEntry 17 }
-- .1.6.1.17

ptpPortDsVlanId OBJECT-TYPE
    SYNTAX      Unsigned32(0..4095)
    MAX-ACCESS   read-write
    STATUS      current

```

DESCRIPTION

"The 802.1Q VLAN Identity to be used by the PTP stack internal to a bridge if portDs.vlanEnable is true; A bridge port would not send a message with a VID of 0 if the port VLAN setting indicates that VLANs are enabled"

REFERENCE "IEC62439-3 A.8.3.4.6; IEEE 802.1Q"

:= { ptpPortDsEntry 18 }

-- .1.6.1.18

ptpPortDsVlanPrio OBJECT-TYPE

SYNTAX Unsigned32(0..7)

MAX-ACCESS read-write

STATUS current

DESCRIPTION

"The 802.1Q priority (0 to 7) to be used by PTP messages originating from the PTP stack of this clock if portDs.vlanEnable is true"

REFERENCE "IEC62439-3 A.8.3.4.7; IEEE 802.1Q"

:= { ptpPortDsEntry 19 }

-- .1.6.1.19

ptpPortDsTwoStepFlag OBJECT-TYPE

SYNTAX TruthValue

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"Enables (if true) 2-step correction at sending of this port (if the port implements it) and indicates that this port uses 2-step correction"

REFERENCE "IEC62439-3 A.8.3.4.8"

:= { ptpPortDsEntry 20 }

-- .1.6.1.20

ptpPortDsPeerIdentity OBJECT-TYPE

SYNTAX ClockIdentity

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The Medium Access Control address of the peer clock in a peer-to-peer delay measurement"

REFERENCE "IEC62439-3 A.8.3.4.9"

:= { ptpPortDsEntry 21 }

-- .1.6.1.21

ptpPortDsPrpAttachement OBJECT-TYPE

SYNTAX PortAttachmentType

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The type of redundant attachment of this port."

REFERENCE "IEC62439-3 A.8.3.4.10"

:= { ptpPortDsEntry 22 }

-- .1.6.1.22

ptpPortDsPrpPairedPort OBJECT-TYPE

SYNTAX Unsigned32(0..4095)

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The portNumber of the other port to which this port is attached, if the port is paired for the purpose of double attachment; a value greater than 4095 indicates the port is not paired"

REFERENCE "IEC62439-3 A.8.3.4.11"

:= { ptpPortDsEntry 23 }

-- .1.6.1.23

ptpPortDsErrorCounter OBJECT-TYPE

SYNTAX Counter32

MAX-ACCESS read-only

STATUS current

DESCRIPTION

"The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that should be normally received given the port state"

REFERENCE "IEC62439-3 A.8.3.4.12"
 ::= { ptpPortDsEntry 24 }
 -- .1.6.1.24

ptpPortDsPeerDelayLim OBJECT-TYPE
 SYNTAX PtpTimeInterval
 UNITS "scaledNanoseconds"
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The value of the peer delay that, if exceeded, will cause an event to be raised"

REFERENCE "IEC62439-3 A.8.3.4.13"
 DEFVAL {'640000'H}
 ::= { ptpPortDsEntry 25 }
 -- .1.6.1.25

ptpTc OBJECT IDENTIFIER ::= { ptpObjects 2 }

ptpTcDefaultDs OBJECT IDENTIFIER ::= { ptpTc 1 }

ptpTcDefaultDsClkIdentity OBJECT-TYPE
 SYNTAX ClockIdentity
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The clockIdentity attribute of the local clock."

REFERENCE "IEC61588 8.3.2.2.1, 7.6.2.1"
 ::= { ptpTcDefaultDs 1 }
 -- .2.1.2

ptpTcDefaultDsNumberPorts OBJECT-TYPE
 SYNTAX Unsigned32(1..4095)
 MAX-ACCESS read-only
 STATUS current
 DESCRIPTION

"The number of PTP ports of the clock (1 if OC), 2 in IEC 62439-3 if Doubly Attached OC and the port is paired"

REFERENCE "IEC61588 8.3.2.2.2"
 ::= { ptpTcDefaultDs 2 }
 -- .2.1.3

ptpTcDefaultDsDelayMech OBJECT-TYPE
 SYNTAX DelayMechanism
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The mechanism used to measure the link delay.

E2E according to Table 9 if the transparent clock uses the end-to-end delay measurement;

P2P according to Table 9 if the transparent clock uses the peer-to-peer delay measurement"

REFERENCE "IEC61588 8.3.2.3.1"
 DEFVAL {2}
 ::= { ptpTcDefaultDs 3 }
 -- .2.1.4

ptpTcDefaultDsPriDomain OBJECT-TYPE
 SYNTAX Unsigned32(0..255)
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION

"The value of transparentClockDefaultDs.primaryDomain of the primary syntonization domain; The initialization value is 0"

REFERENCE "IEC61588 8.3.2.3.2, 10.1"
 ::= { ptpTcDefaultDs 4 }
 -- .2.1.5

ptpTcDefaultDsProfileSet OBJECT-TYPE

SYNTAX ProfileSet
MAX-ACCESS read-write
STATUS current
DESCRIPTION

"The set of profiles which the clock has the ability to support with the required performance (Profile activation is a port property)"

REFERENCE "IEC62439-3 A.8.4.2.1"

::= { ptpTcDefaultDs 5 }

-- .2.1.6

ptpTcDefaultDsTimeInaccuracy OBJECT-TYPE

SYNTAX PtpTimeInterval
UNITS "scaledNanoseconds"
MAX-ACCESS read-only
STATUS current
DESCRIPTION

"The absolute value of the rated worst-case time error introduced between any two ports"

REFERENCE "IEC62439-3 A.8.4.2.2"

DEFVAL {'7FFFFFFF'H}

::= { ptpTcDefaultDs 6 }

-- .2.1.7

ptpTcPortDsTable OBJECT-TYPE

SYNTAX SEQUENCE OF PtpTcPortDsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION

"The table of port dataset related variables. A value of 1 is used in devices with only one port."

::= { ptpTc 2 }

-- .2.2

ptpTcPortDsEntry OBJECT-TYPE

SYNTAX PtpTcPortDsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION

"A list of objects pertaining to a port."

REFERENCE "IEC61588 8.2.5"

INDEX {
ptpTcPortDsIndex }

::= { ptpTcPortDsTable 1 }

-- .2.2.1

PtpTcPortDsEntry ::= SEQUENCE {

ptpTcPortDsIndex Unsigned32,
ptpTcPortDsClockIdentity ClockIdentity,
ptpTcPortDsPortNumber Unsigned32,
ptpTcPortDsLogMinPdelayReqInterv Integer32,
ptpTcPortDsFaultyFlag TruthValue,
ptpTcPortDsPeerMeanLinkDelay PtpTimeInterval,
ptpTcPortDsPortEnabled TruthValue,
ptpTcPortDsDlyAsymmetry PtpTimeInterval,
ptpTcPortDsTwoStepFlag TruthValue,
ptpTcPortDsPeerIdentity ClockIdentity,
ptpTcPortDsPrpAttachement PortAttachmentType,
ptpTcPortDsPrpPairedPort Unsigned32,
ptpTcPortDsErrorCounter Counter32,
ptpTcPortDsPeerDelayLim PtpTimeInterval
}

ptpTcPortDsIndex OBJECT-TYPE

SYNTAX Unsigned32(0..65535)
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION


```

        "Port index"
        ::= { ptpTcPortDsEntry 1 }
-- .2.2.1.1

ptpTcPortDsClockIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The portIdentity.clockIdentity attribute of the local port"
    REFERENCE    "IEC61588 8.3.3.2.1, 7.5.2.2.1"
    ::= { ptpTcPortDsEntry 2 }
-- .2.2.1.2

ptpTcPortDsPortNumber OBJECT-TYPE
    SYNTAX      Unsigned32(0..65535)
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "The portIdentity.portNumber attribute of the local port"
    REFERENCE    "IEC61588 8.3.3.2.1, 7.5.2.3"
    ::= { ptpTcPortDsEntry 3 }
-- .2.2.1.3

ptpTcPortDsLogMinPdelayReqInterv OBJECT-TYPE
    SYNTAX      Integer32(-128..127)
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION
        "The value of transparentClockPortDs.logMinPdelayReqInterval as the logarithm
to the base 2 of the
minPdelayReqInterval"
    REFERENCE    "IEC61588 8.3.3.3.1, 7.7.2.5"
    ::= { ptpTcPortDsEntry 4 }
-- .2.2.1.4

ptpTcPortDsFaultyFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "If the transparentClockPortDs.faultyFlag is set: true; otherwise: false
(initialization value)"
    REFERENCE    "IEC61588 8.3.3.3.2"
    ::= { ptpTcPortDsEntry 5 }
-- .2.2.1.5

ptpTcPortDsPeerMeanLinkDelay OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS        "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS       current
    DESCRIPTION
        "If the transparentClockDefaultDs.delayMechanism is P2P, the value of
transparentClockPortDs.peerMeanPathDelay as an estimate of the current one-way
propagation
delay on the link attached to this port computed using the peer delay mechanism;
otherwise, if the value of the transparentClockDefaultDs.delayMechanism member is
E2E, the value
is zero. The data type should be TimeInterval."
    REFERENCE    "IEC61588 8.3.3.3.3, 11.4"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpTcPortDsEntry 6 }
-- .2.2.1.6

ptpTcPortDsPortEnabled OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-write
    STATUS       current
    DESCRIPTION

```

```

    "Enables the port for PTP traffic and indicates the enabled status of the
port if true"
    REFERENCE      "IEC62439-3 A.8.4.3.1"
    DEFVAL         {1}
    ::= { ptpTcPortDsEntry 7 }
-- .2.2.1.7

ptpTcPortDsDlyAsymmetry OBJECT-TYPE
    SYNTAX      PtpTimeInterval
    UNITS       "scaledNanoseconds"
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The link delay asymmetry to correct the link delay"
    REFERENCE    "IEC61588 7.4.2, IEC62439-3 A.8.4.3.2"
    DEFVAL       {'7FFFFFFF'H}
    ::= { ptpTcPortDsEntry 8 }
-- .2.2.1.8

ptpTcPortDsTwoStepFlag OBJECT-TYPE
    SYNTAX      TruthValue
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "Enables (if true) 2-step correction at sending of this port (if the port
implements it) and indicates that this port uses 2-step correction"
    REFERENCE    "IEC62439-3 A.8.4.3.3"
    DEFVAL       {1}
    ::= { ptpTcPortDsEntry 9 }
-- .2.2.1.9

ptpTcPortDsPeerIdentity OBJECT-TYPE
    SYNTAX      ClockIdentity
    MAX-ACCESS   read-only
    STATUS      current
    DESCRIPTION
        "The Medium Access Control address of the peer clock in a peer-to-peer delay
measurement"
    REFERENCE    "IEC62439-3 A.8.4.3.4"
    ::= { ptpTcPortDsEntry 10 }
-- .2.2.1.10

ptpTcPortDsPrpAttachement OBJECT-TYPE
    SYNTAX      PortAttachmentType
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The type of redundant attachment of this port"
    REFERENCE    "IEC62439-3 A.8.4.3.5"
    DEFVAL       {3}
    ::= { ptpTcPortDsEntry 11 }
-- .2.2.1.11

ptpTcPortDsPrpPairedPort OBJECT-TYPE
    SYNTAX      Unsigned32(0..4095)
    MAX-ACCESS   read-write
    STATUS      current
    DESCRIPTION
        "The portNumber of the other port to which this port is attached, if the port
is paired for the purpose of double attachment; a value greater than 4095 indicates
the port is not paired"
    REFERENCE    "IEC62439-3 A.8.4.3.6"
    DEFVAL       {0}
    ::= { ptpTcPortDsEntry 12 }
-- .2.2.1.12

ptpTcPortDsErrorCounter OBJECT-TYPE
    SYNTAX      Counter32
    MAX-ACCESS   read-only
    STATUS      current

```

DESCRIPTION
 "The error counter that is incremented for each reception error of PTP messages (e.g. time-out), only for those messages that should be normally received given the port state."
 REFERENCE "IEC62439-3 A.8.4.3.7"
 ::= { ptpTcPortDsEntry 13 }
 -- .2.2.1.13

ptpTcPortDsPeerDelayLim OBJECT-TYPE
 SYNTAX PtpTimeInterval
 UNITS "scaledNanoseconds"
 MAX-ACCESS read-write
 STATUS current
 DESCRIPTION
 "The value of the peer delay that, if exceeded, will cause an event to be raised."
 REFERENCE "IEC62439-3 A.8.4.3.8"
 DEFVAL {'640000'H}
 ::= { ptpTcPortDsEntry 14 }
 -- .2.2.1.14

ptpTLVs OBJECT IDENTIFIER ::= { ptpObjects 3 }

ptpAlternateTimeOffsetIndicator OBJECT IDENTIFIER ::= { ptpTLVs 1 }

ptpAtoiTlvTable OBJECT-TYPE
 SYNTAX SEQUENCE OF PtpAtoiTlvEntry
 MAX-ACCESS not-accessible
 STATUS current
 DESCRIPTION
 "Table of ALTERNATE_TIME_OFFSET_INDICATORS"
 REFERENCE "IEC61588 16.3"
 ::= { ptpAlternateTimeOffsetIndicator 1 }
 -- .3.1.1

ptpAtoiTlvEntry OBJECT-TYPE
 SYNTAX PtpAtoiTlvEntry
 MAX-ACCESS not-accessible
 STATUS current
 DESCRIPTION
 "Table entry"
 REFERENCE "IEC61588 16.3"
 INDEX {
 ptpAtoiTlvIndex,
 ptpAtoiTlvTxRxIndex }
 ::= { ptpAtoiTlvTable 1 }
 -- .3.1.1.1

PtpAtoiTlvEntry ::= SEQUENCE {
 ptpAtoiTlvIndex Unsigned32,
 ptpAtoiTlvTxRxIndex TlvTxRxIndex,
 ptpAtoiCurrentOffset Integer32,
 ptpAtoiJumpSeconds Integer32,
 ptpAtoiTimeOfNextJump Integer48,
 ptpAtoiDisplayName OCTET STRING
 }

ptpAtoiTlvIndex OBJECT-TYPE
 SYNTAX Unsigned32(0..65535)
 MAX-ACCESS not-accessible
 STATUS current
 DESCRIPTION
 "Table Index for each alternate time offset indicator"
 REFERENCE "IEC61588 16.3"
 ::= { ptpAtoiTlvEntry 1 }
 -- .3.1.1.1.1

ptpAtoiTlvTxRxIndex OBJECT-TYPE
 SYNTAX TlvTxRxIndex

```

MAX-ACCESS    not-accessible
STATUS        current
DESCRIPTION
    "Table Index for received or sent objects"
REFERENCE     "IEC61588 16.3"
::= { ptpAtoiTlvEntry 2 }
-- .3.1.1.1.2

ptpAtoiCurrentOffset OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    UNITS       "seconds"
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION
        "The value of currentOffset (offset of the alternate time, in seconds, from
the time of this node). The alternate time is the sum of this value and the time of
the node, which is TAI if the PTP time scale is used."
    REFERENCE   "IEC61588 16.3.3.4"
    ::= { ptpAtoiTlvEntry 3 }
-- .3.1.1.1.3

ptpAtoiJumpSeconds OBJECT-TYPE
    SYNTAX      Integer32(-32768..32767)
    UNITS       "seconds"
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION
        "The value of jumpSeconds (size of the next discontinuity, in seconds, of the
alternate time). A value of zero indicates that no discontinuity is expected. A
positive value indicates that the discontinuity
will cause the currentOffset of the alternate time to increase."
    REFERENCE   "IEC61588 16.3.3.5"
    ::= { ptpAtoiTlvEntry 4 }
-- .3.1.1.1.4

ptpAtoiTimeOfNextJump OBJECT-TYPE
    SYNTAX      Integer48
    UNITS       "seconds"
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION
        "The value of timeOfNextJump (value of the seconds portion of the time of the
transmitting node at the time that the next discontinuity will occur). The
discontinuity occurs at the start of the second indicated
by the value of timeOfNextJump."
    REFERENCE   "IEC61588 16.3.3.6"
    ::= { ptpAtoiTlvEntry 5 }
-- .3.1.1.1.5

ptpAtoiDisplayName OBJECT-TYPE
    SYNTAX      OCTET STRING (SIZE (10))
    MAX-ACCESS  read-write
    STATUS      current
    DESCRIPTION
        "The value of displayName (text name of the alternate timescale in ASCII-8).
It represent the time zone, e.g. CET;
a leading character of * indicates that daylight saving time is in effect."
    REFERENCE   "IEC61588 16.3.3.7"
    ::= { ptpAtoiTlvEntry 6 }
-- .3.1.1.1.6

=====
-- Conformance Information
=====

ptpConformance OBJECT IDENTIFIER ::= { ptp 2}
ptpGroups OBJECT IDENTIFIER ::= { ptpConformance 1}
ptpDefaultDsGrp OBJECT-GROUP
    OBJECTS {
        ptpDefaultDsTwoStepFlag,

```

```

    ptpDefaultDsClockIdentity,
    ptpDefaultDsNumberPorts,
    ptpDefaultDsClockClass,
    ptpDefaultDsClockAccuracy,
    ptpDefaultDsOffsetScaledLogVaria,
    ptpDefaultDsPriority1,
    ptpDefaultDsPriority2,
    ptpDefaultDsDomainNumber,
    ptpDefaultDsSlaveOnly,
    ptpDefaultDsProfileSet,
    ptpDefaultDsTimeInaccuracy,
    ptpDefaultDsOffsetFromMasterLim
  }
  STATUS current
  DESCRIPTION
    "Objects in the default dataset"
  ::= {ptpGroups 1}

ptpCurrentDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpCurrentDsStepsRemoved,
    ptpCurrentDsOffsetFromMaster,
    ptpCurrentDsMeanLinkDelay
  }
  STATUS current
  DESCRIPTION
    "Objects in the current dataset"
  ::= {ptpGroups 2}

ptpParentDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpParentDsClockIdentity,
    ptpParentDsPortNumber,
    ptpParentDsStats,
    ptpParentDsObsOfstScdLVar,
    ptpParentDsObsPhChgRate,
    ptpParentDsGmIdentity,
    ptpParentDsGmClass,
    ptpParentDsGmAccuracy,
    ptpParentDsGmOfstScdLVar,
    ptpParentDsGmPriority1,
    ptpParentDsGmPriority2
  }
  STATUS current
  DESCRIPTION
    "Objects in the parent dataset"
  ::= {ptpGroups 3}

ptpTimePropDsGrp  OBJECT-GROUP
  OBJECTS {
    ptpTimePropDsCurrentUtcOffset,
    ptpTimePropDsCurrentUtcOffsetVld,
    ptpTimePropDsLeap59,
    ptpTimePropDsLeap61,
    ptpTimePropDsTimeTraceable,
    ptpTimePropDsFrequencyTraceable,
    ptpTimePropDsPtpTimescale,
    ptpTimePropDsTimeSource
  }
  STATUS current
  DESCRIPTION
    "Objects in the time properties dataset"
  ::= {ptpGroups 4}

ptpForeignMastersDsGrp  OBJECT-GROUP
  OBJECTS {

```

```

    ptpForeignMasterDsPortIdentity,
    ptpForeignMasterAnnounceMessages
}
STATUS current
DESCRIPTION
    "Objects in the foreign master dataset"
::= {ptpGroups 5}

```

```

ptpPortDsGrp OBJECT-GROUP
OBJECTS {
    ptpPortDsClockIdentity,
    ptpPortDsPortNumber,
    ptpPortDsPortState,
    ptpPortDsLogMinDelayReqInterval,
    ptpPortDsPeerMeanLinkDelay,
    ptpPortDsLogAnnounceInterval,
    ptpPortDsAnnounceReceiptTimeout,
    ptpPortDsLogSyncInterval,
    ptpPortDsDelayMechanism,
    ptpPortDsLogMinPdelayReqInterval,
    ptpPortDsVersionNumber,
    ptpPortDsNetworkProtocol,
    ptpPortDsPortEnabled,
    ptpPortDsDlyAsymmetry,
    ptpPortDsProfileId,
    ptpPortDsVlanEnable,
    ptpPortDsVlanId,
    ptpPortDsVlanPrio,
    ptpPortDsTwoStepFlag,
    ptpPortDsPeerIdentity,
    ptpPortDsPrpAttachement,
    ptpPortDsPrpPairedPort,
    ptpPortDsErrorCounter,
    ptpPortDsPeerDelayLim
}
STATUS current
DESCRIPTION
    "Objects in the port dataset"
::= {ptpGroups 6}

```

```

ptpTcDefaultDsGrp OBJECT-GROUP
OBJECTS {
    ptpTcDefaultDsClkIdentity,
    ptpTcDefaultDsNumberPorts,
    ptpTcDefaultDsDelayMech,
    ptpTcDefaultDsPriDomain,
    ptpTcDefaultDsProfileSet,
    ptpTcDefaultDsTimeInaccuracy
}
STATUS current
DESCRIPTION
    "Objects in the transparent clock default dataset"
::= {ptpGroups 7}

```

```

ptpTcPortDsGrp OBJECT-GROUP
OBJECTS {
    ptpTcPortDsClockIdentity,
    ptpTcPortDsPortNumber,
    ptpTcPortDsLogMinPdelayReqInterv,
    ptpTcPortDsFaultyFlag,
    ptpTcPortDsPeerMeanLinkDelay,
    ptpTcPortDsPortEnabled,
    ptpTcPortDsDlyAsymmetry,
    ptpTcPortDsTwoStepFlag,
    ptpTcPortDsPeerIdentity,
    ptpTcPortDsPrpAttachement,
    ptpTcPortDsPrpPairedPort,

```

```

        ptpTcPortDsErrorCounter,
        ptpTcPortDsPeerDelayLim
    }
    STATUS current
    DESCRIPTION
        "Objects in the transparent clock port dataset"
    ::= { ptpGroups 8}

ptpTlvGrp  OBJECT-GROUP
    OBJECTS {
        ptpAtoiCurrentOffset,
        ptpAtoiJumpSeconds,
        ptpAtoiTimeOfNextJump,
        ptpAtoiDisplayName
    }
    STATUS current
    DESCRIPTION
        "Objects in the TLV dataset"
    ::= { ptpGroups 9}

ptpNotificationGrp  NOTIFICATION-GROUP
    NOTIFICATIONS {
        eventFaultyState,
        eventRedundancyLoss,
        eventPeerDelayExceeded
    }
    STATUS current
    DESCRIPTION
        "Notifications"
    ::= { ptpGroups 10}

=====
-- MIB module Compliance statements
=====

ptpCompliances OBJECT IDENTIFIER ::= { ptpConformance 2 }
ptpCompliance MODULE-COMPLIANCE
    STATUS current
    DESCRIPTION
        "Compliance statement for support by of the IEC 61439-8 module"
    MODULE
    MANDATORY-GROUPS {
        eventPeerDelayExceeded,
        ptpDefaultDsGrp,
        ptpCurrentDsGrp,
        ptpParentDsGrp,
        ptpTimePropDsGrp,
        ptpForeignMastersDsGrp,
        ptpPortDsGrp,
        ptpTcDefaultDsGrp,
        ptpTcPortDsGrp,
        ptpTlvGrp,
        ptpNotificationGrp
    }
    ::= { ptpCompliances 1 }
END

```

Bibliography

IEC 61850 (all parts), *Communication networks and systems for power utility automation*

IEC 61850-8-1, *Communication networks and systems for power utility automation – Part 8-1: Specific communication service mapping (SCSM) – Mappings to MMS (ISO 9506-1 and ISO 9506-2) and to ISO/IEC 8802-3*

IEC 61850-9-2, *Communication networks and systems for power utility automation – Part 9-2: Specific communication service mapping (SCSM) – Sampled values over ISO/IEC 8802-3*

IEC 62439-2, *Industrial communication networks – High availability automation networks – Part 2: Media Redundancy Protocol (MRP)*

IEC 62439-6, *Industrial communication networks – High availability automation networks – Part 6: Distributed Redundancy Protocol (DRP)*

IEC 62439-7, *Industrial communication networks – High availability automation networks – Part 7: Ring-based Redundancy Protocol (RRP)*

ISO/IEC 8802-2, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 2: Logical Link Control*

ISO/IEC 9646-7, *Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 7: Implementation Conformance Statements*

RFC 4330 (2006), *The Internet Society, Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

SOMMAIRE

AVANT-PROPOS.....	176
INTRODUCTION.....	178
0.1 Généralités	178
0.2 Modifications par rapport à l'édition précédente	178
0.3 Déclaration de brevet	178
1 Domaine d'application.....	180
2 Références normatives	180
3 Termes, définitions, abréviations, acronymes et conventions	181
3.1 Termes et définitions	181
3.2 Abréviations et acronymes	182
3.3 Conventions.....	183
4 Protocole de redondance en parallèle (PRP)	183
4.1 Principe de fonctionnement du protocole PRP	183
4.1.1 Topologie du réseau PRP.....	183
4.1.2 LAN PRP avec topologie linéaire ou en bus	184
4.1.3 LAN PRP avec topologie en anneau	184
4.1.4 Structure des nœuds DANP	185
4.1.5 Association PRP de nœuds à une seule association	187
4.1.6 Compatibilité entre les nœuds à une seule association et à double association	187
4.1.7 Gestion du réseau.....	187
4.1.8 Implication sur l'application	187
4.1.9 Transition vers des réseaux non redondants	188
4.1.10 Gestion des doublons.....	188
4.1.11 Supervision du réseau.....	194
4.1.12 Interface de gestion de la redondance	194
4.2 Spécifications du protocole PRP	194
4.2.1 Lignes directrices d'installation, de configuration et de réparation	194
4.2.2 Adresses MAC monodiffusion.....	195
4.2.3 Adresses MAC multidiffusion.....	195
4.2.4 Adresses IP	195
4.2.5 Nœuds.....	195
4.2.6 Mode d'acceptation des doublons (essai seulement).....	196
4.2.7 Mode de rejet des doublons	196
4.3 Trame PRP_Supervision	200
4.3.1 Format de la trame PRP_Supervision	200
4.3.2 Contenu de la trame PRP_Supervision	203
4.3.3 Trame PRP_Supervision pour RedBox.....	203
4.3.4 Réception d'une trame PRP_Supervision et NodesTable.....	204
4.4 Nœud de pontage	204
4.5 Constantes	204
4.6 Spécification de service PRP.....	204
5 Redondance transparente haute disponibilité (HSR).....	205
5.1 Objectifs du protocole HSR	205
5.2 Principe de fonctionnement du protocole HSR	205
5.2.1 Fonctionnement de base avec une topologie en anneau	205

5.2.2	Structure des nœuds DANH	208
5.2.3	Topologie.....	209
5.2.4	Structure d'une RedBox	219
5.3	Spécifications du nœud HSR	220
5.3.1	Fonctionnement HSR	220
5.3.2	Réception de données par le DANH depuis son interface de couche de liaison.....	221
5.3.3	Réception de données par le DANH depuis un port HSR.....	221
5.3.4	Règles de transfert du DANH	222
5.3.5	CoS	224
5.3.6	Synchronisation de l'horloge.....	224
5.3.7	Accès déterministe au support.....	224
5.4	Spécifications de la RedBox HSR	224
5.4.1	Propriétés de la RedBox.....	224
5.4.2	Réception RedBox via l'interconnexion	225
5.4.3	Transfert RedBox sur l'anneau	226
5.4.4	Réception RedBox via un port HSR	226
5.4.5	Réception RedBox via son interface de couche de liaison	228
5.4.6	Gestion de ProxyNodeTable de la RedBox	228
5.4.7	CoS RedBox	228
5.4.8	Synchronisation de l'horloge RedBox.....	228
5.4.9	Accès au support RedBox	228
5.5	Spécification de la QuadBox.....	229
5.6	Méthode de rejet des doublons.....	229
5.7	Format de trame pour HSR.....	229
5.7.1	Format de trame pour toutes les trames	229
5.7.2	Trame HSR_Supervision.....	231
5.8	Constantes	233
5.9	Spécification de service HSR.....	234
6	Enoncé de conformité de mise en œuvre de protocole (PICS).....	236
7	Base d'informations de gestion (MIB) PRP/HSR	237
Annexe A (normative) Synchronisation d'horloges sur des chemins redondants dans l'IEC 62439-3.....		254
A.1	Vue d'ensemble	254
A.2	Association à des LAN redondants par une horloge frontière.....	254
A.3	Association à des LAN redondants par des horloges ordinaires à double association.....	255
A.4	Mapping PRP vers PTP	258
A.4.1	Scénarios et rôles d'appareils.....	258
A.4.2	Fonctionnement dans PRP	259
A.4.3	Spécification de la configuration	261
A.4.4	Spécifications d'un DANP en tant que DAC.....	262
A.4.5	Modèle d'horloge d'une RedBox pour PTP	262
A.5	Mapping HSR vers PTP.....	282
A.5.1	Trafic PTP dans HSR	282
A.5.2	Spécifications des nœuds HSR	285
A.5.3	Horloges redondantes dans HSR.....	286
A.5.4	Association d'une MC à un LAN externe	286
A.6	Mapping PRP vers HSR	287

A.6.1	Méthodes de connexion	287
A.6.2	Connexion PRP-HSR par une BC	287
A.6.3	Connexion PRP-HSR par des TC	288
A.7	Modèle d'horloge à double association	289
A.7.1	Diagramme d'états	289
A.7.2	Supervision du port	292
A.7.3	BMCA pour les ports couplés	293
A.7.4	Sélection de l'état des ports	295
A.8	Ensembles de données PTP pour la haute disponibilité	295
A.8.1	Généralités	295
A.8.2	Types de données	295
A.8.3	Ensembles de données pour les horloges ordinaires ou frontières	296
A.8.4	Objet pour les horloges transparentes	300
Annexe B (normative) Profil PTP pour l'automatisation des systèmes électriques – Association d'une horloge redondante		303
B.1	Domaine d'application	303
B.2	Spécification du profil PTP	303
B.3	Association d'une horloge redondante	303
Annexe C (normative) Profils PTP pour les réseaux d'automatisme à haute disponibilité		304
C.1	Domaine d'application	304
C.2	Spécification du profil PTP	304
C.3	Types d'horloges	304
C.4	Informations communes sur la spécification du protocole	305
C.5	Spécification du protocole pour le profil d'automatisation L3E2E	305
C.6	Spécification du protocole pour le profil d'automatisation L2P2P	305
C.7	Exigences de temporisation	306
C.7.1	Conditions de mesure	306
C.7.2	Imprécision temporelle du réseau	306
C.7.3	Éléments du réseau	306
C.7.4	Exigences pour les horloges grands maîtres	306
C.7.5	Exigences pour les TC	307
C.7.6	Exigences pour les BC	307
C.7.7	Exigences pour les convertisseurs de support	307
C.7.8	Exigences pour les liaisons	308
C.8	Ingénierie de réseau	308
C.9	Réglages par défaut	309
C.10	Traitement des horloges redondantes	310
C.11	Enoncé de conformité de mise en œuvre de protocole (PICS)	310
C.11.1	Conventions	310
C.11.2	PICS	310
Annexe D (informative) Tutoriel relatif au protocole PTP pour l'IEC 62439-3		312
D.1	Objectif	312
D.2	Exactitude et précision	312
D.3	Types d'horloges PTP	313
D.4	Options PTP principales	316
D.5	Communication de couches 2 et 3	316
D.6	Correction en 1 étape/en 2 étapes	316
D.6.1	Correction du temps dans les TC	316

D.6.2	Conversion 2 étapes/1 étape	318
D.7	Mesure du retard de ligne entre extrémités	321
D.7.1	Méthode générale	321
D.7.2	Mesure du retard de ligne entre extrémités avec correction d'horloge en 1 étape	321
D.7.3	Mesure du retard de ligne entre extrémités avec correction d'horloge en 2 étapes	323
D.7.4	Calcul de retard de ligne entre extrémités à partir des messages Delay_Req/Delay_Resp	324
D.8	Calcul de retard de ligne entre homologues	325
D.8.1	Calcul de retard de ligne entre homologues avec correction d'horloge en 1 étape	325
D.8.2	Calcul de retard de ligne entre homologues avec correction d'horloge en 2 étapes	327
Annexe E (normative) Base d'informations de gestion pour les horloges à une seule association et les horloges à double association		329
Bibliographie		354
Figure 1 – Exemple PRP de réseau redondant général		183
Figure 2 – Exemple PRP de réseau redondant constitué de deux LAN (topologie en bus) 184		
Figure 3 – Exemple PRP d'anneau redondant avec des SAN et des DANP		185
Figure 4 – PRP avec deux DANP qui communiquent		186
Figure 5 – RedBox PRP, transition d'un LAN simple vers un LAN double		188
Figure 6 – Trame PRP étendue par une RCT		189
Figure 7 – Trame PRP à étiquette VLAN étendue par une RCT		190
Figure 8 – Trame remplie PRP fermée par une RCT		191
Figure 9 – Limites de l'algorithme de rejet de doublons		192
Figure 10 – Exemple HSR d'une configuration en anneau pour le trafic multidiffusion		206
Figure 11 – Exemple HSR d'une configuration en anneau pour le trafic monodiffusion		207
Figure 12 – Structure HSR d'un DANH		209
Figure 13 – Exemple HSR de topologie qui utilise deux réseaux indépendants		210
Figure 14 – Exemple HSR du couplage par homologues de deux anneaux		211
Figure 15 – Exemple HSR d'anneaux connectés		212
Figure 16 – Exemple HSR du couplage de deux réseaux locaux PRP redondants à un anneau		213
Figure 17 – Exemple HSR de couplage du nœud d'un anneau à des LAN PRP redondants		215
Figure 18 – Exemple HSR de couplage d'un anneau à deux LAN PRP		216
Figure 19 – Exemple HSR de couplage de trois anneaux à un LAN PRP		217
Figure 20 – Exemple HSR de topologie en maille		218
Figure 21 – Structure HSR d'une RedBox		219
Figure 22 – Trame HSR sans balise VLAN		229
Figure 23 – Trame HSR avec balise VLAN		230
Figure 24 – Nœud HSR avec compteurs de gestion		235
Figure 25 – RedBox HSR avec compteurs de gestion		236

Figure A.1 – Horloge à double association comme BC (MCA est la meilleure horloge maître)	254
Figure A.2 – Horloge à double association quand MCA est la meilleure horloge maître.....	256
Figure A.3 – Horloges à double association quand OC1 est la meilleure horloge maître	257
Figure A.4 – Eléments des réseaux PRP	259
Figure A.5 – Connexion d'une horloge maître à une horloge ordinaire via PRP.....	261
Figure A.6 – RedBox PRP en tant que BC (OC3 et BC7 sont les meilleurs maîtres).....	263
Figure A.7 – Modèle d'horloge DABC RedBox	264
Figure A.8 – RedBox PRP en tant que DABC avec E2E – BC7 est l'horloge maître	266
Figure A.9 – RedBox PRP en tant que DABC avec E2E – temporisation.....	267
Figure A.10 – RedBox PRP en tant que DABC avec P2P – OC5 est la meilleure horloge maître	268
Figure A.11 – RedBox PRP en tant que DABC avec P2P – temporisation.....	269
Figure A.12 – RedBox PRP en tant que DATC avec E2E – flux de signaux.....	271
Figure A.13 – RedBox PRP en tant que DATC avec E2E – temporisation.....	273
Figure A.14 – RedBox PRP en tant que DATC avec P2P	275
Figure A.15 – RedBox PRP en tant que DATC avec P2P – temporisation	276
Figure A.16 – RedBox PRP en tant que SLTC avec E2E.....	279
Figure A.17 – RedBox PRP en tant que SLTC avec E2E – temporisation	280
Figure A.18 – RedBox PRP en tant que SLTC avec P2P	281
Figure A.19 – HSR avec une GMC	283
Figure A.20 – Messages PTP envoyés et reçus par un nœud HSR (1 étape).....	284
Figure A.21 – Messages PTP envoyés et reçus par un nœud HSR (2 étapes)	285
Figure A.22 – Association d'une GMC à un anneau HSR par l'intermédiaire d'une RedBox en tant que TC.....	287
Figure A.23 – Couplage PRP à HSR par des BC.....	288
Figure A.24 – Couplage PRP à HSR par des TC.....	289
Figure A.25 – Etats des ports avec les transitions pour le fonctionnement redondant.....	290
Figure A.26 – BMCA pour les maîtres redondants.....	294
Figure D.1 – Exemple de précision et d'exactitude.....	312
Figure D.2 – Principe du protocole PTP.....	313
Figure D.3 – Eléments du protocole PTP	316
Figure D.4 – Retards et logique d'horodatage dans les TC.....	317
Figure D.5 – Correction du message Sync en 1 étape/2 étapes (entre homologues)	318
Figure D.6 – Conversion 2 étapes/1 étape dans les TC.....	320
Figure D.7 – Conversion 2 étapes/1 étape – Vue des messages	321
Figure D.8 – Mesure du retard de ligne entre extrémités avec correction d'horloge en 1 étape.....	322
Figure D.9 – Mesure du retard de ligne entre extrémités avec correction d'horloge en 2 étapes	324
Figure D.10 – Mesure du retard de ligne entre homologues avec correction d'horloge en 1 étape	326
Figure D.11 – Mesure du retard de ligne entre homologues avec correction d'horloge en 2 étapes	328

Tableau 1 – Cas de rejet des doublons.....	193
Tableau 2 – Ensemble de données de surveillance.....	196
Tableau 3 – Attributs de NodesTable.....	197
Tableau 4 – Trame PRP_Supervision sans balise VLAN.....	201
Tableau 5 – Trame PRP_Supervision avec balise VLAN (facultative).....	202
Tableau 6 – Contenu de la trame PRP_Supervision.....	203
Tableau 7 – TLV PRP_Supervision pour RedBox.....	203
Tableau 8 – Constantes PRP.....	204
Tableau 9 – Trame HSR_Supervision sans balise VLAN.....	231
Tableau 10 – Trame HSR_Supervision avec balise VLAN facultative.....	232
Tableau 11 – Constantes HSR.....	234
Tableau A.1 – Etats.....	291
Tableau A.2 – Transitions.....	292
Tableau A.3 – Variables.....	292
Tableau C.1 – Attributs PTP pour le profil d'automatisation industrielle.....	309
Tableau C.2 – PICS pour les horloges.....	311

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 RLV

COMMISSION ÉLECTROTECHNIQUE INTERNATIONALE

**RÉSEAUX DE COMMUNICATION INDUSTRIELS –
RÉSEAUX D'AUTOMATISME A HAUTE DISPONIBILITE –****Partie 3: Protocole de redondance en parallèle (PRP)
et redondance transparente de haute disponibilité (HSR)**

AVANT-PROPOS

- 1) La Commission Electrotechnique Internationale (IEC) est une organisation mondiale de normalisation composée de l'ensemble des comités électrotechniques nationaux (Comités nationaux de l'IEC). L'IEC a pour objet de favoriser la coopération internationale pour toutes les questions de normalisation dans les domaines de l'électricité et de l'électronique. A cet effet, l'IEC – entre autres activités – publie des Normes internationales, des Spécifications techniques, des Rapports techniques, des Spécifications accessibles au public (PAS) et des Guides (ci-après dénommés "Publication(s) de l'IEC"). Leur élaboration est confiée à des comités d'études, aux travaux desquels tout Comité national intéressé par le sujet traité peut participer. Les organisations internationales, gouvernementales et non gouvernementales, en liaison avec l'IEC, participent également aux travaux. L'IEC collabore étroitement avec l'Organisation Internationale de Normalisation (ISO), selon des conditions fixées par accord entre les deux organisations.
- 2) Les décisions ou accords officiels de l'IEC concernant les questions techniques représentent, dans la mesure du possible, un accord international sur les sujets étudiés, étant donné que les Comités nationaux de l'IEC intéressés sont représentés dans chaque comité d'études.
- 3) Les Publications de l'IEC se présentent sous la forme de recommandations internationales et sont agréées comme telles par les Comités nationaux de l'IEC. Tous les efforts raisonnables sont entrepris afin que l'IEC s'assure de l'exactitude du contenu technique de ses publications; l'IEC ne peut pas être tenue responsable de l'éventuelle mauvaise utilisation ou interprétation qui en est faite par un quelconque utilisateur final.
- 4) Dans le but d'encourager l'uniformité internationale, les Comités nationaux de l'IEC s'engagent, dans toute la mesure possible, à appliquer de façon transparente les Publications de l'IEC dans leurs publications nationales et régionales. Toutes divergences entre toutes Publications de l'IEC et toutes publications nationales ou régionales correspondantes doivent être indiquées en termes clairs dans ces dernières.
- 5) L'IEC elle-même ne fournit aucune attestation de conformité. Des organismes de certification indépendants fournissent des services d'évaluation de conformité et, dans certains secteurs, accèdent aux marques de conformité de l'IEC. L'IEC n'est responsable d'aucun des services effectués par les organismes de certification indépendants.
- 6) Tous les utilisateurs doivent s'assurer qu'ils sont en possession de la dernière édition de cette publication.
- 7) Aucune responsabilité ne doit être imputée à l'IEC, à ses administrateurs, employés, auxiliaires ou mandataires, y compris ses experts particuliers et les membres de ses comités d'études et des Comités nationaux de l'IEC, pour tout préjudice causé en cas de dommages corporels et matériels, ou de tout autre dommage de quelque nature que ce soit, directe ou indirecte, ou pour supporter les coûts (y compris les frais de justice) et les dépenses découlant de la publication ou de l'utilisation de cette Publication de l'IEC ou de toute autre Publication de l'IEC, ou au crédit qui lui est accordé.
- 8) L'attention est attirée sur les références normatives citées dans cette publication. L'utilisation de publications référencées est obligatoire pour une application correcte de la présente publication.

La Norme internationale IEC 62439-3 a été établie par le sous-comité 65C: Réseaux industriels, du comité d'études 65 de l'IEC: Mesure, commande et automation dans les processus industriels.

Cette troisième édition annule et remplace la deuxième édition parue en 2012. Cette édition constitue une révision technique.

Cette édition inclut les modifications techniques majeures suivantes par rapport à l'édition précédente:

- a) corrections techniques et extension des spécifications;
- b) prise en compte de la synchronisation des horloges selon l'IEC 61588 lors de la mesure du délai entre extrémités et la mesure existante du délai entre homologues dans le protocole PRP.

Le texte de cette norme est issu des documents suivants:

FDIS	Rapport de vote
65C/834/FDIS	65C/841/RVD

Le rapport de vote indiqué dans le tableau ci-dessus donne toute information sur le vote ayant abouti à l'approbation de cette norme.

Cette publication a été rédigée selon les Directives ISO/IEC, Partie 2.

Cette Norme internationale doit être lue conjointement avec l'IEC 62439-1.

Une liste de toutes les parties de la série IEC 62439, publiées sous le titre général *Réseaux industriels de communication – Réseaux d'automatisme à haute disponibilité*, peut être consultée sur le site web de l'IEC.

Le comité a décidé que le contenu de cette publication ne sera pas modifié avant la date de stabilité indiquée sur le site web de l'IEC sous "<http://webstore.iec.ch>" dans les données relatives à la publication recherchée. A cette date, la publication sera

- reconduite,
- supprimée,
- remplacée par une édition révisée, ou
- amendée.

IMPORTANT – Le logo "colour inside" qui se trouve sur la page de couverture de cette publication indique qu'elle contient des couleurs qui sont considérées comme utiles à une bonne compréhension de son contenu. Les utilisateurs devraient, par conséquent, imprimer cette publication en utilisant une imprimante couleur.

INTRODUCTION

0.1 Généralités

L'IEC 62439-3 fait partie de la série IEC 62439 *Réseaux industriels de communication – Réseaux d'automatisme à haute disponibilité*. Elle spécifie les protocoles de redondance transparente PRP et HSR. Elle a été adoptée par l'IEC TC57 WG10 pour servir de méthode de redondance pour les réseaux d'automatisme de la sous-station demandeuse qui fonctionne sur des réseaux de couche 2, conformément à l'IEC 61850-8-1 et à l'IEC 61850-9-2.

Le principe de redondance transparente a été étendu aux horloges qui fonctionnent conformément au protocole PTP (Precision Time Protocol), comme défini dans l'IEC 61588 et associé aux réseaux redondants. Deux variantes sont spécifiées: L3E2E pour les horloges qui fonctionnent sur des réseaux de couche 3 avec une mesure du retard de ligne entre extrémités (E2E) et L2P2P pour les horloges qui fonctionnent sur des réseaux de couche 2 avec une mesure du retard de ligne entre homologues (P2P).

0.2 Modifications par rapport à l'édition précédente

Les principales modifications par rapport à l'IEC 62439-3:2012 sont:

- Le paragraphe 4.1.10.3 a été réécrit pour expliquer la façon de calculer le rejet des doublons pour différentes vitesses.
- L'Annexe A a été revue pour servir de conception générale pour des horloges à double association utilisables pour une mesure du retard de ligne entre extrémités (E2E) et entre homologues (P2P); le principe du fonctionnement de ports couplés a été désormais spécifié en termes de diagramme d'états basés sur l'IEC 61588:2009.
- L'Annexe B de l'IEC 62439-3:2012 a été supprimée; ses propriétés sont mentionnées en 5.3.7.
- La nouvelle Annexe B rend obligatoire la prise en charge de la redondance pour l'IEC/IEEE 61850-9-3 qui spécifie les horloges à double association sur la couche 2 avec une mesure du retard de ligne entre homologues.
- L'Annexe C spécifie deux profils d'une horloge de précision pour l'automatisation industrielle:
L3E2E pour la couche 3, mesure du retard entre extrémités et
L2P2P pour la couche 2, mesure du retard entre homologues.
- L'Annexe D contient les informations de tutoriel relatives à l'IEC 61588:2009 qui sont nécessaires pour comprendre les annexes ci-dessus. Ces informations se trouvaient à l'Annexe A de l'IEC 62439-3:2012.
- L'Annexe E (MIB) comprend la base d'informations de gestion SNMP à utiliser pour des horloges à une seule association et à double association, quel que soit le profil.

0.3 Déclaration de brevet

La Commission Electrotechnique Internationale (IEC) attire l'attention sur le fait qu'il est déclaré que la conformité avec les dispositions du présent document peut impliquer l'utilisation d'un brevet intéressant le filtrage de trames redondantes dans le nœud d'un réseau (Siemens Aktiengesellschaft – EP 2127329, US 8184650, CN 101611615B) traité en 5.2.3.3.

L'IEC ne prend pas position quant à la preuve, à la validité et à la portée de ces droits de propriété.

Le détenteur de ces droits de propriété a donné l'assurance à l'IEC qu'il consent à négocier des licences avec des demandeurs du monde entier, à des termes et conditions raisonnables et non discriminatoires. A ce propos, l'énoncé du détenteur des droits de propriété est enregistré à l'IEC. Des informations peuvent être demandées à:

Siemens Aktiengesellschaft
Oto-Hahn-Ring 6
81379 Munich, Allemagne

La Commission Electrotechnique Internationale (IEC) attire l'attention sur le fait qu'il est déclaré que la conformité avec les dispositions du présent document peut impliquer l'utilisation de brevets sur la réception de trames redondantes et non redondantes (ABB Research Ltd – EP 1825657, US 8582426, CN 101057483, IN 254425) traitée en 4.2.7, sur l'identification d'un mauvais câblage des appareils (ABB Technology AG – EP 2163024, US 8344736, CN 101689985) traitée en 4.3, sur l'appareil critique avec disponibilité accrue (ABB Research Ltd – EP 2090950) traité en 4.4 et sur les nœuds de couplage d'un anneau pour les réseaux à haute disponibilité (ABB Research Ltd – US 8582424, EP 2327185, CN 102106121) traités en 5.2.3.

L'IEC ne prend pas position quant à la preuve, à la validité et à la portée de ces droits de propriété.

Le détenteur de ces droits de propriété a donné l'assurance à l'IEC qu'il consent à négocier des licences avec des demandeurs du monde entier, à des termes et conditions raisonnables et non discriminatoires. A ce propos, l'énoncé du détenteur des droits de propriété est enregistré à l'IEC. Des informations peuvent être demandées à:

ABB Schweiz AG
Intellectual Property CH-IP (CH-150016-L)
Brown Boveri Strasse 6
CH-5400 Baden, Suisse
ch-ip.patent@abb.com

L'attention est d'autre part attirée sur le fait que certains des éléments du présent document peuvent faire l'objet de droits de propriété autres que ceux qui ont été mentionnés ci-dessus. L'IEC ne saurait être tenue pour responsable de ne pas avoir identifié de tels droits de brevets et de ne pas avoir signalé leur existence.

L'ISO (www.iso.org/patents) et l'IEC (<http://patents.iec.ch>) maintiennent à disposition des bases de données en ligne des brevets relatifs à leurs normes. Les utilisateurs sont encouragés à consulter ces bases de données pour obtenir l'information la plus récente sur les droits de propriété.

RÉSEAUX DE COMMUNICATION INDUSTRIELS – RÉSEAUX D'AUTOMATISME A HAUTE DISPONIBILITE –

Partie 3: Protocole de redondance en parallèle (PRP) et Redondance transparente haute disponibilité (HSR)

1 Domaine d'application

La série IEC 62439 s'applique aux réseaux d'automatisme à haute disponibilité basés sur la technologie Ethernet.

Cette partie de l'IEC 62439 spécifie deux protocoles de redondance conçus pour assurer une reprise transparente en cas de défaillance unique de liaison interpont ou de pont au sein du réseau, basés sur le schéma: transmission en parallèle d'informations dupliquées.

2 Références normatives

Les documents suivants sont cités en référence de manière normative, en intégralité ou en partie, dans le présent document et sont indispensables pour son application. Pour les références datées, seule l'édition citée s'applique. Pour les références non datées, la dernière édition du document de référence s'applique (y compris les éventuels amendements).

IEC 60050-191, *Vocabulaire Electrotechnique International – Chapitre 191: Sécurité de fonctionnement et qualité de service*

IEC 61588:2009, *Precision clock synchronization protocol for networked measurement and control systems* (disponible en anglais seulement)

IEC TR 61850-90-4, *Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines* (disponible en anglais seulement)

IEC 62439-1, *Réseaux industriels de communication – Réseaux d'automatisme à haute disponibilité – Partie 1: Concepts généraux et méthodes de calcul*

IEC/IEEE 61850-9-3:—, *Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation (proposed IEC 61850-9-3)*, (disponible en anglais seulement)¹

ISO/IEC/IEEE 8802-3:2014, *Standard for Ethernet* (disponible en anglais seulement)

IEEE 802.1D:2004, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges* (disponible en anglais seulement)

IEEE 802.1Q:2014, *IEEE Standard for Local and metropolitan area networks – Media Access Control (MAC) Bridges and Virtual Bridge Local Area Network* (disponible en anglais seulement)

¹ A paraître.

IETF RFC 2578, *Structure of Management Information Version 2 (SMIv2)* (disponible en anglais seulement)

IETF RFC 3418, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)* (disponible en anglais seulement)

3 Termes, définitions, abréviations, acronymes et conventions

3.1 Termes et définitions

Pour les besoins du présent document, les termes et définitions donnés dans l'IEC 60050-191 et dans l'IEC 62439-1, ainsi que les suivants, s'appliquent.

3.1.1

rôle grand maître

horloge ordinaire ou horloge frontière capable de jouer le rôle de grand maître

Note 1 à l'article: Une horloge à rôle grand maître n'est pas forcément reliée à une source d'heure reconnue.

3.1.2

interconnexion

liaison entre deux hiérarchies de réseaux

3.1.3

imprécision temporelle du réseau

imprécision temporelle évaluée ou estimée entre une horloge esclave et le signal de référence temporel des horloges grands maîtres, en fonction de la pire topologie réseau

3.1.4

RedBox

appareil qui permet de connecter des nœuds à une seule association à un réseau redondant

3.1.5

QuadBox

appareil à quatre ports qui permet de connecter deux anneaux HSR homologues, agit comme un nœud HSR dans chaque anneau et peut filtrer et transmettre le trafic d'anneau en anneau

3.1.6

trame HSR

trame qui transporte EtherType en tant que HSR_ethertype

3.1.7

logique de commutation

matériel qui transmet une trame d'un port à un autre et peut fournir la commutation à la volée

3.1.8

erreur de temporisation

écart par rapport à la référence temporelle utilisée pour mesurer la synchronisation, évaluée sur un court intervalle de temps

3.1.9

imprécision temporelle

erreur de temporisation qui n'est pas dépassée par 99,7 % des mesures, évaluée sur une série de 1 000 mesures (environ 20 min) en régime établi

Note 1 à l'article: Dans le présent document, les mots "ponts" et "pontage" sont respectivement synonymes des mots "commutateurs" et "commutation" lorsqu'ils s'appliquent à une connectivité de couche 2.

3.2 Abréviations et acronymes

Pour les besoins du présent document, les abréviations et acronymes donnés dans l'IEC 62439-1, ainsi que les suivants, s'appliquent.

ARP	Address Resolution Protocol (protocole de résolution d'adresse)
BC	Boundary Clock (horloge frontière, comme défini dans l'IEC 61588:2009)
BMCA	Best Master Clock Algorithm, algorithme de la meilleure horloge maître, comme défini dans l'IEC 61588:2009
BOC	Boundary Ordinary Clock (horloge ordinaire frontière)
CoS	Class of Service (classe de service (qualité de service de couche 2))
DAC	Doubly Attached Clock (horloge à double association)
DABC	Doubly Attached Boundary Clock (horloge frontière à double association)
DAN	Doubly Attached Node (nœud à double association)
DANH	Doubly Attached Node implementing HSR (nœud à double association de mise en œuvre du HSR)
DANP	Doubly Attached Node implementing PRP (nœud à double association de mise en œuvre du PRP)
DATC	Doubly Attached Transparent Clock (horloge transparente à double association)
E2E	End-to-End (entre extrémités (mesure du retard))
GMC	Grandmaster Clock (horloge grand maître)
HC	Hybrid Clock (horloge hybride)
HSR	High-availability Seamless Redundancy (redondance transparente haute disponibilité)
IP	Internet Protocol (protocole Internet)
L2P2P	Layer 2 Peer-to-Peer (couche 2 entre homologues)
L3E2E	Layer 3 End-to-End (couche 3 entre extrémités)
LAN	Local Area Network (réseau local)
LLC	Link Layer Control (contrôle de couche de liaison)
LRE	Link Redundancy Entity (entité de redondance de liaison)
MAC	Medium Access Control (contrôle d'accès au support)
MC	Master Clock (horloge maître)
MIB	Management Information Base (base d'informations de gestion)
OC	Ordinary Clock (horloge ordinaire, comme défini dans l'IEC 61588:2009)
P2P	Peer to Peer (entre homologues (mesure du retard))
PICS	Protocol Implementation Conformance Statement (énoncé de conformité de mise en œuvre de protocole)
PRP	Parallel Redundancy Protocol (protocole de redondance en parallèle)
PT	Protocol Type (type de protocole)
RCT	Redundancy Check Trailer (queue de contrôle de redondance)
RSTP	Rapid Spanning Tree Protocol (protocole STA rapide)
SAN	Singly Attached Node (nœud à une seule association)
SAC	Singly Attached Clock (horloge à une seule association)
SLTC	Stateless Transparent Clock (horloge transparente sans état)
SNMP	Simple Network Management Protocol (protocole simple de gestion de réseau)
TC	Transparent Clock (horloge transparente, comme défini dans l'IEC 61588:2009)
TCP	Transport Control Protocol (protocole de contrôle de transport)
UDP	User Data Protocol (protocole de données d'utilisateur)
VDAN	Virtual Doubly Attached Node (nœud à association virtuelle double (SAN vu d'une RedBox))

3.3 Conventions

Cette partie de la série IEC 62439 applique les conventions définies dans l'IEC 62439-1.

4 Protocole de redondance en parallèle (PRP)


4.1 Principe de fonctionnement du protocole PRP

4.1.1 Topologie du réseau PRP

Ce protocole de redondance met en œuvre la redondance dans les nœuds plutôt que dans le réseau (voir les définitions de l'IEC 62439-1), à l'aide de nœuds à double association qui obéissent à PRP (les DANP).

Un DANP est associé à deux réseaux locaux (LAN) indépendants de topologie similaire, nommés LAN_A et LAN_B, qui fonctionnent en parallèle. Un DANP source envoie la même trame sur les deux LAN et un DANP de destination la reçoit des deux LAN dans un délai donné, utilise la première trame et rejette le doublon.

La Figure 1 présente un réseau redondant composé de deux réseaux locaux qui peuvent avoir une topologie quelconque, en arbre, en anneau ou maillée, par exemple.



IEC

Anglais	Français
A-frame	trame A
B-frame	trame B
bridge	pont
bridged local area network (ring)	LAN ponté (anneau)
bridged local area network (tree)	LAN ponté (arborescence)

Figure 1 – Exemple PRP de réseau redondant général


Les protocoles des deux LAN sont identiques au niveau MAC-LLC, mais ils peuvent différer en termes de performances et de topologie. Les retards de transmission peuvent aussi être différents, en particulier si l'un des réseaux se reconfigure, à l'aide du protocole RSTP, pour remédier à une défaillance interne, par exemple.

Les deux LAN suivent des règles de configuration qui permettent le bon fonctionnement des protocoles de gestion de réseau, comme les protocoles ARP et SNMP.

Les deux LAN ne sont pas connectés entre eux et sont considérés comme indépendants en cas de défaillance. La redondance peut être compromise par des points de défaillance uniques, comme une alimentation commune ou une connexion directe dont la défaillance provoque la panne des deux réseaux. Les directives d'installation de ce document donnent à l'installateur des lignes directrices pour assurer l'indépendance en cas de défaillance.

4.1.2 LAN PRP avec topologie linéaire ou en bus

Pour une topologie plus simple, la Figure 2 présente un réseau PRP constitué de deux réseaux locaux avec une topologie linéaire, qui peut aussi être une topologie en bus.



IEC

Figure 2 – Exemple PRP de réseau redondant constitué de deux LAN (topologie en bus)

4.1.3 LAN PRP avec topologie en anneau

Les deux LAN peuvent présenter une topologie en anneau; voir Figure 3.

La redondance dans chaque anneau est assurée par un mécanisme différent, par exemple celui que définit l'IEEE 802.1D ou d'autres protocoles de l'IEC 62439.




Figure 3 – Exemple PRP d'anneau redondant avec des SAN et des DNP

4.1.4 Structure des nœuds DNP

Chaque nœud dispose de deux ports qui fonctionnent en parallèle et sont associés aux mêmes couches supérieures de la pile de communication via l'entité LRE (Link Redundancy Entity, entité de redondance de liaison), comme le montre la Figure 4.




Figure 4 – PRP avec deux DANP qui communiquent

Pour la communication de base, l'entité LRE présente dans ses couches supérieures la même interface qu'un adaptateur réseau non redondant. Les couches supérieures ne reconnaissent donc pas la redondance.

L'entité LRE a deux fonctions: le traitement des doublons et la gestion de la redondance.

Lorsqu'elle reçoit une trame des couches supérieures du nœud, l'entité LRE lui ajoute une queue de contrôle de redondance (RCT) avec un numéro de séquence, puis envoie la trame par ses deux ports pratiquement en même temps. Les deux trames sont pratiquement identiques, sauf leur identificateur LAN (et la somme de contrôle).

Les deux trames transitent par les deux LAN avec des retards différents et arrivent pratiquement au même moment au nœud de destination.

Lorsqu'elle reçoit des trames du réseau, l'entité LRE transmet la première trame reçue d'une paire aux couches supérieures de son nœud et rejette le doublon (s'il arrive). Elle supprime la RCT, si cela est exigé.

Pour gérer la redondance et vérifier la présence d'autres DANP, une entité LRE envoie régulièrement des trames PRP_Supervision et peut évaluer celles envoyées par d'autres DANP.

4.1.5 Association PRP de nœuds à une seule association

Les nœuds à une seule association (SAN) peuvent être associés de deux manières:

Les SAN ne peuvent être associés directement qu'à un seul LAN. Ce type de SAN ne peut communiquer qu'avec d'autres SAN sur le même LAN. Par exemple, à la Figure 1, SAN A1 peut communiquer avec SAN A2, mais pas avec SAN B1 ou SAN B2. Les SAN peuvent communiquer (de manière non redondante) avec tous les DANP.

Les SAN peuvent être associés par l'intermédiaire d'un boîtier de redondance (RedBox) aux deux LAN, comme l'indique la Figure 1 pour SAN R1 et SAN R2 (voir aussi 4.1.9). Ces SAN peuvent communiquer avec tous les DANP et SAN, SAN A1 et SAN R1 peuvent communiquer, par exemple.

NOTE Les SAN ne reconnaissent pas la présence de PRP. Ce sont par exemple des ordinateurs ou des imprimantes disponibles dans le commerce.

Dans certaines applications, seuls les appareils essentiels pour la disponibilité doivent disposer d'une association double, les postes de travail des opérateurs, par exemple, tandis que la majorité des appareils sont des SAN. Un DANP peut tirer parti de l'infrastructure de base du protocole PRP pour être associé à deux ponts différents du même LAN (un anneau, par exemple) et utiliser des protocoles autres que PRP pour reconfigurer le réseau en cas de défaillance. Le DANP se comporte alors comme un pont, conformément à l'IEEE 802.1D. Par exemple, l'élément de pont peut mettre en œuvre le protocole RSTP ou un sous-ensemble de RSTP, où aucun trafic n'est transmis entre les ports. Ces capacités facultatives ne sont pas détaillées dans le présent document. Le mode pris en charge est spécifié à l'Article 6.

4.1.6 Compatibilité entre les nœuds à une seule association et à double association

Les nœuds à une seule association (SAN), par exemple, les ordinateurs portables de maintenance ou les imprimantes d'un LAN, peuvent être connectés à n'importe quel LAN. Les ponts sont toujours des SAN. Un SAN connecté à un LAN ne peut pas communiquer directement avec un SAN connecté à l'autre LAN. Ces SAN ne reconnaissent pas la redondance PRP, mais les DANP génèrent un trafic que ces SAN comprennent.

La condition est toutefois que les SAN ignorent la RCT dans les trames, ce qui est le cas puisqu'un SAN ne peut pas distinguer la RCT du remplissage Ethernet.

Inversement, les DANP comprennent le trafic généré par les SAN, puisque ceux-ci n'ajoutent pas de RCT. Ils ne transmettent qu'une trame à leurs couches supérieures, car le trafic des SAN utilise un seul LAN.

4.1.7 Gestion du réseau

Un nœud possède la même adresse MAC sur les deux ports et un seul ensemble d'adresses IP est affecté à cette adresse. La redondance est donc transparente pour les couches supérieures, ce qui permet en particulier au protocole ARP de fonctionner de la même manière qu'un SAN. Les ponts d'un LAN ne sont pas des appareils à double association. Tous les ponts gérés possèdent donc des adresses IP (et MAC) différentes. Un outil de gestion de réseau est de préférence un DANP qui peut accéder aux nœuds et aux ponts des deux LAN.

4.1.8 Implication sur l'application


Quand la même trame provient de deux ports dont la différence de temps est significative, dans certains cas exceptionnels, une trame pourrait être par erreur rejetée comme doublon. Il existe par conséquent un mécanisme qui élimine les trames au bout d'un certain temps. Ce mécanisme de vieillissement a dans de rares cas pour effet que les deux doublons sont acceptés. La couche de transport (par exemple, TCP) doit alors filtrer le doublon ou

l'application doit tolérer le doublon. Ce cas est normal, puisqu'un LAN ne garantit pas l'absence de doublons.

Le mécanisme de vieillissement prend en compte le taux maximal de répétition des trames et la différence la plus défavorable dans le retard de propagation, comme décrit en 4.1.10.3.

4.1.9 Transition vers des réseaux non redondants

Le mécanisme de rejet des doublons peut être mis en œuvre par une RedBox qui assure la transition entre un SAN et les LAN doublés, comme le montre la Figure 5 La RedBox simule les SAN connectés derrière elle (VDAN ou DAN virtuels) et diffuse en multidiffusion des trames de supervision en leur nom. La RedBox, elle-même un DANP, possède sa propre adresse IP à des fins de gestion, même si elle peut aussi exécuter des fonctions d'application.



Anglais	Français
bridge	pont
non-redundant network	réseau non redondant
singly attached nodes	nœuds à une connexion
local application	application locale
transceivers	émetteurs-récepteurs

Figure 5 – RedBox PRP, transition d'un LAN simple vers un LAN double

4.1.10 Gestion des doublons

4.1.10.1 Mode d'acceptation des doublons (essai seulement)

Le mode d'acceptation des doublons est utilisé à des fins d'essai; il permet de vérifier que les doublons sont écartés par la couche de liaison, pas par les protocoles de couche supérieure.

En mode d'acceptation des doublons, un émetteur est configuré pour envoyer une trame normale vers les deux ports.

En mode d'acceptation des doublons, un récepteur est configuré pour accepter et transférer les deux trames (si elles sont arrivées) à ses couches supérieures.

Un nœud indique son mode (acceptation/rejet des doublons) dans ses trames de supervision.

NOTE Les couches supérieures rejettent aussi les doublons, mais moins efficacement que la couche de liaison.

4.1.10.2 Mode de rejet des doublons

4.1.10.2.1 Principe


Pour permettre aux récepteurs de détecter les doublons, la LRE émettrice ajoute un champ à 6 octets avec un numéro de séquence, RCT, aux 2 trames, comme le montre la Figure 6.

La LRE réceptrice utilise le numéro de séquence de la RCT et l'adresse MAC source pour détecter les doublons. Elle ne transfère que la première trame d'une paire à ses couches supérieures. Pour délester le processeur d'application, la LRE peut être dotée d'un préprocesseur indépendant, d'un contrôleur Ethernet intelligent ou de matériel.

Le fait d'assurer le suivi des doublons permet aussi d'améliorer la supervision de la redondance. Le récepteur peut éventuellement supprimer la RCT, puis transférer la (première) trame à sa couche supérieure.

La RCT est composée des champs suivants:

- numéro de séquence de 16 bits (SeqNr);
- identificateur de LAN de 4 bits (LanId);
- taille de trame de 12 bits (LSDUsize);
- suffixe de 16 bits (PRPsuffix).



IEC

Anglais	Français
octet position	position d'octet
preamble	préambule
LSDU size	taille LSDU
PRP suffix	suffixe PRP
time	temps
frame without redundancy control	trame sans contrôle de redondance
Redundancy Control Trailer	queue de contrôle de redondance

Figure 6 – Trame PRP étendue par une RCT

La LRE envoie les deux trames (presque identiques) sur les deux LAN.

La LRE de réception peut ensuite détecter les doublons en fonction de la RCT.

NOTE 1 Cette méthode considère que des SAN sont aussi présents sur le réseau et que les trames envoyées par les SAN peuvent être rejetées par erreur en tant que doublons, car leur champ de fin affiche le suffixe PRP correct, le même numéro de séquence et la même taille. Les SAN envoient cependant les trames sur un seul LA. Leur adresse MAC source ne sera donc pas identique à celle d'une trame sur l'autre LAN, de sorte qu'une trame issue d'un SAN ne sera jamais rejetée.

NOTE 2 La réception d'une RCT bien formée ne prouve pas que le nœud source est un DANP. Un DANP accepte la trame même si le chemin est incorrect.


4.1.10.2.2 Utilisation de PRPsuffix

PRPsuffix identifie les trames PRP et les distingue des autres protocoles qui ajoutent aussi une queue à leurs données utiles.

4.1.10.2.3 Utilisation de LSDUsize

Le champ LSDUsize de 12 bits contient la longueur de l'unité de données de service de liaison (LSDU) en octets. Il permet au récepteur de distinguer les trames PRP des trames non redondantes aléatoires dans le cadre d'un contrôle supplémentaire.

Dans les VLAN, des balises VLAN de trame peuvent être ajoutées ou supprimées lors du transit via un pont. Pour que le champ de longueur soit indépendant du balisage VLAN, seules la LSDU et la RCT sont prises en compte dans LSDUsize, comme le montre la Figure 7.




IEC

Anglais	Français
octet position	position d'octet
LSDU size	taille LSDU
PRP suffix	suffixe PRP
time	temps
preamble	préambule

Figure 7 – Trame PRP à étiquette VLAN étendue par une RCT

Si le récepteur détecte que la trame est fermée par le suffixe PRP correct, que les 12 bits qui précèdent le suffixe PRP correspondent à la taille LSDU et que l'identificateur de LAN correspond à celui du port destinataire, la trame est candidate au rejet.

Comme les trames courtes doivent disposer d'un remplissage pour atteindre la taille de trame minimale de 64 octets, l'émetteur inclut déjà le remplissage pour éviter d'analyser la trame à partir de la fin, comme le montre la Figure 8 La taille de trame minimale est étendue à 70 octets, car l'entité LRE ou une RedBox peut supprimer la RCT.



Anglais	Français
octet position	position d'octet
LSDU size	taille LSDU
PRP suffix	suffixe PRP
time	temps
preamble	préambule
padding	remplissage

Figure 8 – Trame remplie PRP fermée par une RCT

La taille du remplissage est portée à 74 octets si une balise VLAN est utilisée.

NOTE Une trame à balise VLAN peut traverser plusieurs ponts, ce qui supprime ou insère des balises VLAN. Si l'émetteur respecte la règle Ethernet pour envoyer des trames d'une taille minimale de 68 octets pour une trame à balise VLAN et de 64 octets pour une trame sans balise VLAN, il n'y aura pas de situations dans lesquelles il existe un remplissage avant et après la RCT.

4.1.10.2.4 Utilisation de LanId

Le champ LanId de 4 bits de la RCT transporte un identificateur différent pour LAN_A ou LAN_B, notamment les codes 1010 ("A") et 1011 ("B"). La trame A et la trame B diffèrent donc d'un bit (et au niveau de la FCS). Le récepteur vérifie que la trame provient du LAN approprié. Il ne rejette pas une trame qui provient du mauvais LAN, car il pourrait s'agir d'une trame légitime dont le suffixe PRP et les informations de longueur se trouvent dans les 28 derniers bits, mais il incrémente le compteur d'erreurs CntErrWrongLanA ou CntErrWrongLanB, ce qui indique une erreur de configuration. Ce type d'erreur est permanent. Il est donc rapidement détecté.

4.1.10.2.5 Utilisation de SeqNr

Le champ SeqNr de 16 bits est incrémenté de un pour chaque trame qu'un DANP envoie. Le doublet {adresse MAC source, numéro de séquence} identifie de manière unique les différentes copies d'une même trame.

4.1.10.3 Algorithme de rejet des doublons et configuration

L'algorithme de rejet des doublons n'est pas spécifié. Par exemple, des tables de nœuds, des tables de hachage, le principe de FIFO ou le suivi des numéros de séquence peuvent être utilisés. Quel que soit l'algorithme, il est conçu de manière à ne jamais rejeter une trame légitime alors que l'acceptation occasionnelle d'un doublon peut être tolérée. Comme le champ SeqNr de 16 bits est remis à zéro après 65 536 trames et que les numéros de séquence ne sont pas nécessairement contigus, les entrées qui ont dépassé une durée EntryForgetTime spécifiée sont purgées.

La Figure 9 représente les conditions de rejet ou d'acceptation des trames.

Le temps t_{skew} est la différence de temps entre les arrivées des deux copies de la même trame (du LAN A au LAN B et du LAN B au LAN A) dans un cas particulier.

Le temps maximal $t_{skewMax}$ entre deux copies est une propriété réseau, estimée par le concepteur réseau à partir du nombre de ponts et du trafic pour une application particulière, par exemple 12 ms.

Le temps $t_{aliasMin}$ est l'intervalle minimal possible entre deux trames légitimes différentes avec le même numéro de séquence (après 65 536 incréments) dans une application particulière.

NOTE Un nœud qui se réinitialise avec un numéro de séquence arbitraire pourrait toutefois causer un raccourcissement du temps $t_{aliasMin}$, comme cela sera détaillé ultérieurement.

Dans un réseau 100 Mbit/s, $t_{aliasMin}$ est l'intervalle minimal théorique de répétition de trames de la plus petite trame (64 octets) issue de la même source ($65\,536 \times 6,72 \mu s$) = 440 ms.

Dans un réseau 1 Gbit/s, $t_{aliasMin}$ s'élève à 44 ms.

Cela constitue la situation la plus défavorable, car les applications envoient rarement des trames de taille minimale de manière contiguë.

EXEMPLE: Un appareil IEC 61850-9-2 qui envoie des trames de valeurs échantillonnées à 4,8 kHz présentera un $t_{aliasMin} = 13$ s.




Figure 9 – Limites de l'algorithme de rejet de doublons

En considérant un appareil, les entrées qui figurent dans la liste de doublons demeurent enregistrées durant un temps qui varie entre $t_{resideMin}$ et $t_{resideMax}$. Ces valeurs dépendent de la mise en œuvre, elles sont appropriées pour un certain nombre d'applications et elles peuvent dépendre du trafic si la taille de la liste est limitée. Il est recommandé que les valeurs $t_{resideMin}$ et $t_{resideMax}$ soient réglables selon le débit de données.

Les cas suivants de la Figure 9 sont considérés dans le Tableau 1:

Tableau 1 – Cas de rejet des doublons

Cas	Condition	Description
Cas 1	$t_{\text{skew}} > t_{\text{resideMin}}$	Le doublon est toujours rejeté. Cette situation est souhaitable.
Cas 2	$t_{\text{resideMin}} < t_{\text{skew1}} < t_{\text{resideMax}}$	Le doublon n'est parfois pas rejeté si l'entrée a été purgée. Cette situation est tolérée si elle reste occasionnelle.
Cas 3	$t_{\text{skew1}} > t_{\text{resideMax}}$	Le doublon n'est jamais rejeté, il est traité comme une nouvelle trame et il crée une nouvelle entrée dans la liste des doublons. Cette situation ne se produit jamais si le réseau a été correctement dimensionné avec $(t_{\text{skewMax}} < t_{\text{resideMax}})$.
Cas 4	$t_{\text{skew2}} > t_{\text{resideMax}}$	Le doublon n'est jamais rejeté, mais une trame légitime est rejetée, car l'entrée créée par le LAN B est en conflit avec l'alias du LAN A. Cette situation ne se produit jamais si les nœuds sont correctement conçus.
Cas 5	$t_{\text{skew2}} > t_{\text{aliasMin}}$	Le doublon est rejeté pour la mauvaise raison, mais cela ne compromet pas l'exploitation. Cette situation ne se produit jamais si les nœuds ont été correctement conçus.

Les conditions d'un fonctionnement correct sont les suivantes:

- Pour l'ingénieur réseau:

$$t_{\text{skewMax}} < (t_{\text{aliasMin}})/2 \text{ (pas d'alias)}$$

- Pour l'ingénieur appareil:

$$t_{\text{resideMin}} > t_{\text{skewMax}} \text{ (reconnaît tous les doublons) et}$$

$$t_{\text{resideMax}} < (t_{\text{aliasMin}})/2 \text{ (empêche les rejets non souhaités)}$$

En pratique, $t_{\text{resideMin}}$ est souvent donné par le nombre de trames conservées dans le tampon et dépend donc du trafic, tandis que $t_{\text{resideMax}}$ est géré par un temporisateur. Cette valeur est réglable par le biais du paramètre de gestion `IreDupListResideMaxTime` (voir Article 7)

Un nœud qui se réinitialise pourrait générer des alias à un intervalle inférieur à t_{aliasMin} , car le numéro de séquence au redémarrage est arbitraire. Pour permettre au temps de l'ensemble des nœuds de destination possibles de vider toutes les entrées du nœud de la liste des doublons, un nœud peut ne pas se réinitialiser plus vite que `NodeRebootInterval` (voir Tableau 8). `NodeRebootInterval` vaut pour tous les nœuds du réseau; il est donc indépendant de la vitesse de liaison.

Tous les nœuds vident donc leurs entrées avant `NodeRebootInterval`. Le temps de résidence maximal d'une entrée quelle que soit la vitesse (et la valeur de $t_{\text{resideMax}}$) est `EntryForgetTime`.

La valeur par défaut d'`EntryForgetTime` dans le Tableau 8 (400 ms) permet en principe un t_{skewMax} de 800 ms, en prenant pour hypothèse un t_{aliasMin} de 1 600 ms. Cela convient à toutes les applications envisagées.

Pour satisfaire aux équations ci-dessus, $t_{\text{resideMax}}$ est généralement beaucoup plus petit que `EntryForgetTime`, notamment dans les réseaux 1 Gbit/s.

4.1.10.4 Restriction de taille des trames

L'ajout de la RCT peut générer des trames surdimensionnées qui dépassent la valeur `maxValidSize` prévue par l'ISO/IEC/IEEE 8802-3:2014 (1 522 octets).

Il est donc prévu que chaque composant de réseau des LAN prenne en charge les trames surdimensionnées jusqu'au `maxValidSize` prévu par l'ISO/IEC/IEEE 8802-3:2014 plus 6 octets de la RCT (1 528 octets).

NOTE Beaucoup de ponts sont dimensionnés pour des trames à double balise VLAN (Q-in-Q) d'une taille maximale de 1 526 octets. La plupart des contrôleurs Ethernet sont certifiés jusqu'à 1 528 octets. Les ponts transfèrent correctement les trames jusqu'à 1 536 octets.

4.1.11 Supervision du réseau

L'état de chaque LAN et des appareils associés (nœuds et ponts) est surveillé, sans quoi la redondance a peu d'utilité.

Le récepteur vérifie que les trames sont correctement reçues sur les deux canaux. Il gère les compteurs d'erreurs que la gestion de réseau peut lire.

A cette fin, les émetteurs et les récepteurs peuvent gérer des tables qui enregistrent l'heure de réception de la dernière trame en provenance d'un autre nœud, l'heure d'envoi d'une trame de multidiffusion ou de diffusion et d'autres informations de protocole.

La supervision repose sur chaque DANP qui envoie régulièrement une trame PRP_Supervision qui permet de vérifier l'intégrité du réseau et la présence des nœuds. Dans le même temps, ces trames permettent de vérifier quels appareils sont des DANP, les adresses MAC qu'ils utilisent et le mode d'exploitation qu'ils prennent en charge: accepter les doublons ou rejeter les doublons.

4.1.12 Interface de gestion de la redondance

Les appareils et les liaisons redondants sont inutiles en l'absence d'une gestion de réseau qui supervise la redondance et qui exige des actions de maintenance.

La LRE présente une interface de gestion de réseau qui permet de suivre l'état de chaque LAN, en particulier de détecter rapidement les défaillances lorsque le taux d'erreur augmente. A cette fin, la LRE gère pour chaque port un compteur des messages reçus et des messages reçus avec une erreur.

Les états du LAN apparaissent sous forme d'objets SNMPv1 ou SNMPv2/v3, ce qui permet aux mêmes outils de gérer les nœuds et les ponts.

4.2 Spécifications du protocole PRP

4.2.1 Lignes directrices d'installation, de configuration et de réparation

4.2.1.1 Généralités

Les lignes directrices suivantes s'appliquent lors de l'installation. Elles ne s'appliquent pas aux essais de conformité des appareils.

4.2.1.2 Disposition des LAN

Le réseau doit être composé de deux LAN avec des propriétés similaires: chacun peut transporter le trafic qui existerait en l'absence de redondance.

Le DANP doit être associé aux deux LAN.

Les SAN qui ont besoin de communiquer entre eux doivent être associés au même LAN ou aux deux LAN via une RedBox.

4.2.1.3 Etiquetage des câbles

Les deux LAN doivent utiliser des câbles clairement identifiés A et B.

Si des couleurs sont utilisées pour identifier les câbles, il convient d'utiliser la couleur rouge pour A et la couleur bleue pour B.

NOTE L'objectif est de permettre aux personnes ne percevant pas les couleurs de les distinguer.

4.2.1.4 Identification des ponts

Il convient que les ponts des deux LAN aient une étiquette ou une couleur distincte pour A et B.

4.2.1.5 Fonctionnement indépendant

La disposition des deux LAN doit être adaptée à l'hypothèse d'indépendance en cas de défaillance.

4.2.1.6 Configuration

Tous les DANP doivent être configurés avec la même valeur LifeCheckInterval (voir 4.5).

4.2.2 Adresses MAC monodiffusion

Les ports A et B d'un DANP doivent être configurés avec la même adresse MAC.

Cette adresse MAC doit être unique au sein du réseau.

Les SAN connectés à un seul LAN ne doivent pas avoir la même adresse MAC qu'un autre nœud du réseau global (LAN_A plus LAN_B).

4.2.3 Adresses MAC multidiffusion

Tous les nœuds du réseau doivent être configurés de manière à utiliser la même adresse multidiffusion pour la supervision du réseau (voir 4.3).

4.2.4 Adresses IP

Les adresses IP de chaque nœud ou de chaque pont doivent être uniques dans tout le réseau (LAN_A plus LAN_B).

NOTE 1 L'unicité s'applique à toutes les adresses IP d'un appareil.

Un DANP doit avoir les mêmes adresses IP vu du LAN_A ou vu du LAN_B.

Les ponts sur le LAN_A et le LAN_B sont considérés comme des SAN et doivent avoir des adresses IP différentes pour la gestion du réseau, même s'ils occupent le même emplacement.

NOTE 2 Un schéma possible consiste à distinguer les adresses IP d'un bit, c'est-à-dire à allouer les adresses IP paires aux appareils LAN_A et les adresses IP impaires aux appareils LAN_B, par exemple.

4.2.5 Nœuds

4.2.5.1 Types de nœuds

Les nœuds à double association conformes au protocole de redondance en parallèle (DANP) doivent avoir deux ports (port A et port B) dotés des mêmes capacités; ils pourraient notamment être utilisés seuls si un seul LAN était connecté, puisque le port A est connecté à LAN_A et le port B à LAN_B.

Les nœuds à une seule association (SAN) ne disposent que d'un seul port pour ce protocole. Aucune exigence particulière ne s'applique à eux.

4.2.5.2 Etiquetage des connecteurs

Le Paragraphe 4.2.5.2 s'applique à un DANP qui utilise deux LAN de nature similaire.

Les connecteurs de chaque LAN doivent être étiquetés explicitement comme A et B.

Lorsque les connecteurs sont placés verticalement, il convient que le LAN_A soit le connecteur supérieur et que le LAN_B soit le connecteur inférieur en position normale de l'appareil.

Lorsque les connecteurs sont placés horizontalement, il convient que le connecteur de gauche soit le LAN_A et que le connecteur de droite soit le LAN_B, depuis le côté où les câbles ou les fibres sont branchés.

Les connecteurs redondants doivent pouvoir être retirés et mis en place indépendamment.

4.2.6 Mode d'acceptation des doublons (essai seulement)

4.2.6.1 Envoi

Un émetteur en mode d'acceptation des doublons doit envoyer la trame qu'il reçoit de ses couches supérieures sans la modifier sur ses deux ports, afin que les deux trames apparaissent sur les LAN respectifs.

4.2.6.2 Réception

Un récepteur en mode d'acceptation des doublons doit transférer les trames reçues d'un port vers ses couches supérieures.

NOTE L'acceptation des doublons est un réglage d'essai qui n'est pas utilisé pour l'exploitation.

4.2.7 Mode de rejet des doublons

4.2.7.1 Ensemble de données de surveillance

Un nœud doit gérer un ensemble de données de surveillance constitué des éléments du Tableau 2.

Tableau 2 – Ensemble de données de surveillance

Nom d'objet	Caractère	Description
SendSeq	obligatoire	Numéro de séquence de 16 bits utilisé par ce nœud pour l'envoi
CntErrorA et CntErrorB	facultatif	Pour chaque port A et B, compteur de 32 bits indiquant le nombre de trames erronées reçues
CntReceivedA et CntReceivedB	facultatif	Pour chaque port A et B, compteur de 32 bits indiquant le nombre de trames reçues
CntErrWrongLanA et CntErrWrongLanB	facultatif	Pour chaque port A et B, compteur de 32 bits indiquant le nombre de discordances
ActiveA et ActiveB	facultatif	Pour chaque port A et B, statut de la liaison de communication
NodesTable	facultatif	Table de tous les autres nœuds que ce nœud détecte

4.2.7.2 NodesTable

Un nœud peut gérer un NodesTable pour le besoin de la supervision du réseau. Si un NodesTable est utilisé, il doit contenir les éléments donnés au Tableau 3 pour chaque entrée.

Tableau 3 – Attributs de NodesTable

Argument	Définition	Type de données
MacAddress	Adresse MAC du nœud source distant (6 octets)	OctetString6
CntReceivedA	Nombre de trames reçues de ce nœud sur LAN_A	Unsigned32
CntReceivedB	Nombre de trames reçues de ce nœud sur LAN_B	Unsigned32
CntErrWrongLanA	Nombre de trames reçues de ce nœud avec un identificateur LAN incorrect sur LAN_A	Unsigned32
CntErrWrongLanB	Nombre de trames reçues de ce nœud avec un identificateur LAN incorrect sur LAN_B	Unsigned32
TimeLastSeenA	Heure de réception de la dernière trame depuis ce nœud sur LAN_A	TimeTicks
TimeLastSeenB	Heure de réception de la dernière trame depuis ce nœud sur LAN_B	TimeTicks
SanA	True si le nœud distant est très probablement un SAN accessible sur le port A	Boolean1
SanB	True si le nœud distant est très probablement un SAN accessible sur le port B	Boolean1

Un nœud doit supprimer une entrée NodesTable d'un autre nœud lorsque le délai écoulé depuis la réception d'une trame de ce nœud sur TimeLastSeenA et TimeLastSeenB dépasse NodeForgetTime (voir 4.5).

NOTE 1 NodesTable est rempli par les trames reçues et permet de distinguer les SAN des DANP. Le tableau NodesTable permet à un nœud d'identifier un autre nœud comme un SAN sur un seul LAN. Dans ce cas, le nœud pourrait envoyer des trames sans RCT vers le port de ce LAN seulement. Un nœud sans NodesTable envoie toujours toutes les trames avec une RCT sur deux ports, même si la destination est un SAN. Dans des applications présentant principalement du trafic de diffusion, l'impact sur le réseau est limité.

NOTE 2 L'attribut clé de NodesTable est MacAddress, comme reçu dans la trame PRP_Supervision envoyée par un DANP ou une trame reçue d'un SAN.

NOTE 3 La plupart de ces attributs existent non seulement en une instance par nœud physique distant, mais aussi en tant qu'instances distinctes pour chaque adresse de diffusion/multidiffusion utilisée par ce nœud. Dans certains cas, ils existent aussi pour chaque adresse de diffusion/multidiffusion utilisée par ce nœud (local).

NOTE 4 TimeTicks est l'unité de temps locale, pas un temps global.

4.2.7.3 Queue de contrôle de redondance (RCT)

La RCT (Redundancy Control Trailer) ajoutée aux trames DANP doit être composée de six octets, structurés de la manière suivante (dans l'ordre de transmission):

- un numéro de séquence de 16 bits (SeqNr) transmis avec les 8 bits de poids fort du premier octet qui reflète le compteur SendSeq (voir 4.2.7.1);
- un identificateur de LAN de 4 bits (LanId) transmis comme les 4 bits de poids fort du troisième octet, qui transporte respectivement la séquence 1010 pour LAN_A et 1011 pour LAN_B;
- une taille de LSDU de 12 bits (LSDUsize) dont les 4 bits de poids fort sont transmis dans les 4 bits de poids faible du troisième octet et dont les 8 bits de poids faible sont transmis dans le quatrième octet, ce qui indique la taille en octets de la LSDU entre la fin du champ Protocol Type (PT), comme défini dans l'ISO/IEC/IEEE 8802-3:2014 and IEEE 802.1Q (décalage d'octets 12-13 sans en-tête LAN ou 16-17 avec en-tête VLAN) et la RCT, à l'exclusion du PT, ainsi que le remplissage possible et la RCT elle-même;
- un PRPsuffix de 16 bits, défini sur 0x88FB.

NOTE 1 Le remplissage inséré avant la RCT est inclus dans la taille de LSDU.

NOTE 2 L'objectif est de faire en sorte que LSDUsize soit identique dans PRP et HSR.

4.2.7.4 Envoi (mode de rejet des doublons)

4.2.7.4.1 Envoi

Lorsqu'il envoie une frame issue de ses couches supérieures, un nœud doit appliquer la procédure ci-dessous:

If the destination node is not registered as a SAN in the NodesTable or if the node has no NodesTable	
	If the frame carries a IEEE 802.1Q tag and has a size smaller than 64 octets (without FCS)
	pad the frame to 64 octets as specified in ISO/IEC/IEEE 8802-3:2014
	If the frame does not carry a IEEE 802.1Q tag and has a size smaller than 60 octets (without FCS)
	pad the frame to 60 octets as specified in ISO/IEC/IEEE 8802-3:2014
	increment SendSeq (see 4.2.7.1), wrapping over through 0; append the 6-octet RCT with a LanId "A" for portA and LanID "B" for portB; append the 4-octet FCS; send the frame with LanId "A" through portA and the frame with LanId "B" through portB;
Else	
	send the frame to the port(s) where the SAN is registered;

NOTE 1 Des nœuds simples ne disposent pas de NodesTable et n'enregistrent pas les nœuds distants comme DANP ou SAN. Ils ajoutent toujours la RCT et l'envoient sur les deux ports.

NOTE 2 Le remplissage total à 70 octets (sans balise VLAN) considère que la LRE ou une RedBox pourrait supprimer la RCT à 6 octets. Le remplissage à 74 octets (avec balise VLAN) considère en outre qu'un port d'extrémité pourrait supprimer la balise de priorité à 4 octets. De cette manière, la frame donnée à la destination respecte la taille de frame minimale de 64 octets et la RCT est toujours placée juste avant la FCS.

4.2.7.4.2 Temporisation

Il convient que le décalage entre l'envoi des deux frames d'une paire soit réduit au minimum pour réduire l'asymétrie entre les retards de transmission.

En particulier, il convient que les frames en attente d'envoi soient retirées de la file d'attente d'envoi si le port correspondant n'est pas opérationnel.

4.2.7.5 Récepteur (mode de rejet des doublons)

4.2.7.5.1 Réception

Lors de la réception d'une frame sur un port, le récepteur doit appliquer la procédure ci-dessous:

If the frame is erroneous,	
	increment the error counter of the respective port CntErrorsA or CntErrorsB (see 4.2.7.1)
	ignore the frame;
Else if the destination MAC address belongs to the reserved addresses in IEEE 802.1Q-2004 Table 7-10,	
	treat the frame at the link layer;
Else if the frame is a duplicate (see 4.2.7.5.2 and 4.2.7.5.3),	
	discard the frame;
Else	
	forward the frame to the upper layers.

Les trames dont l'adresse MAC de destination est conforme à l'IEEE 802.1Q:2014, Tableau 8-1, Tableau 8-2 et Tableau 8-3 sont traitées au niveau de la couche de liaison et ne sont pas supposées transporter de RCT, même en cas d'envoi par un DANP.

4.2.7.5.2 Identification des trames associées au mode de rejet des doublons

Un récepteur doit identifier un doublon potentiel comme une trame valide:

- dont les 16 derniers bits avant la somme de contrôle correspondent au PRPsuffix de 16 bits,
- dont les 12 bits avant le PRPsuffix correspondent à la taille physique de la LSDU (à l'exclusion de ProtocolType PT, mais avec la RCT) et
- dont le LanId correspond à l'identité LAN du port de réception, 1010 (A) ou 1011 (B).

4.2.7.5.3 Méthode de rejet des doublons

La méthode de rejet des doublons ne doit pas rejeter une trame envoyée seule ou les deux trames d'une paire. Exceptionnellement, les deux trames d'une paire peuvent être transmises aux couches de protocole supérieures.

La méthode de rejet des doublons doit oublier une entrée après EntryForgetTime (voir Tableau 8).

4.2.7.5.4 Réception transparente

Un récepteur configuré pour TransparentReception (par défaut) ne doit pas supprimer la RCT avant de transférer la trame à ses couches supérieures.

Un récepteur non configuré pour TransparentReception doit supprimer la RCT sur les trames s'il a détecté la présence de la RCT.

4.2.7.5.5 Réception et NodesTable

Si le nœud prend en charge un NodesTable, il doit, à réception d'une trame, appliquer la procédure ci-dessous:

If this frame is not a PRP_Supervision frame or a HSR Supervision frame and its source MAC address is not yet in the NodesTable,		
		create an entry in the NodesTable for that source MAC address assuming it is a SanA or a SanB, depending which LAN the frame arrives on;
Else if this frame is a PRP_Supervision frame		
	If its source is not yet in the NodesTable	
		create an entry in the NodesTable for that source assuming DANP Duplicate Accept or Duplicate Discard according to the PRP_Supervision frame contents;
	Else	
		update the source status to DANP Duplicate Accept or Duplicate Discard;
record the local time at which the frame was received in the TimeLastSeenA, respectively TimeLastSeenB fields of the NodesTable for that source;		
increment by one (wrapping through 0) the counters CntReceivedA, respectively CntReceivedB of the NodesTable for that source.		
If it receives the frame from LAN_B from a node registered as SanA, or over LAN_A from a node registered as SanB,		
		set SanA = SanB = 1 for that source;
If the LanId does not match the identity of the port,		
		increment the CntErrWrongLanA, respective CntErrWrongLanB counter of the source device.

NOTE 1 La mise à jour de SanA et SanB permet de déplacer un SAN du LAN_A vers le LAN_B et inversement. Dans ce cas, le DANP enverra les données sur les deux LAN et, après NodeForgetTime, limitera l'envoi au LAN approprié.

NOTE 2 La réception d'une RCT bien formée n'est pas un critère suffisant pour déclarer sa source en tant que DANP, certains protocoles répondant avec une trame identique à celle reçue.

4.3 Trame PRP_Supervision

4.3.1 Format de la trame PRP_Supervision

Chaque DANP doit diffuser en multidiffusion une trame PRP_Supervision sur chacun de ses deux ports au format spécifié au Tableau 4 (pas de VLAN) ou au Tableau 5 (avec VLAN) respectivement, à une fréquence déterminée par LifeCheckInterval (voir 4.5).

Tableau 4 – Trame PRP_Supervision sans balise VLAN

0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0					msb	U/L	I/G								
2	PRP_DestinationAddress = multidiffusion (01-15-4E-00-01-XX)														
4								lsb							
6					msb	U/L	0								
8	SourceAddress (adresse MAC du nœud)														
10								lsb							
12	SupEtherType = 0x88FB														
14	SupPath					SupVersion < 64									
16	SupSequenceNumber														
18	TLV1.Type = 20 ou 21							TLV1.Length = 6							
20					msb	U/L	0								
22	MacAddress (adresse MAC du DANP)														
24								lsb							
26	TLV2.Type = 30							TLV2.Length = 6							
28					msb	U/L	0								
30	RedBoxMacAddress														
32								lsb							
34	TLV0.Type = 0							TLV0.Length = 0							
36															
	Remplissage à 70 octets (pas de VLAN)														
60	SeqNr														
62	LanId (0x1010 ou 0x1011)					LSDUsize = 52									
64	PRPsuffix = 0x88FB														
66	FCS														
68															
70															

Tableau 5 – Trame PRP_Supervision avec balise VLAN (facultative)

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14																
16																
18																
20																
22																
24						msb	U/L	0								
26																
28									lsb							
30																
32						msb	U/L	0								
34																
36									lsb							
38																
40																
42																
44																
46																
48																
50																
52																
54																
56																
58																
60																
62																
64																
66																
68																
70																
72																
74																

NOTE 1 Les octets avec un décalage de 14 à 17 ne sont insérés que si une balise VLAN conforme à la norme IEEE 802.1Q est utilisée.

NOTE 2 La taille de la trame est respectivement de 70 octets (sans balise VLAN) ou 74 octets (avec balise VLAN) afin d'éviter le remplissage si une entité supprime la balise VLAN ou la RCT.

NOTE 3 SeqNr, LanId et LSDUsize sont réutilisés lorsque la trame est transmise à un réseau HSR. SupEtherType est identique au suffixe PRP. Il est aussi utilisé pour HSR.

NOTE 4 La longueur des adresses MAC est, par hypothèse, de 6 octets.

4.3.2 Contenu de la trame PRP_Supervision

Le contenu d'une trame PRP_Supervision est défini dans le Tableau 6.

Tableau 6 – Contenu de la trame PRP_Supervision

Nom de champ	Description
PRP_DestinationAddress	Adresse multidiffusion 01-15-4E-00-01-XX réservée pour le protocole PRP. Par défaut, XX est "00", mais, en cas de conflits, XX peut être configuré pour prendre n'importe quelle valeur comprise entre 0x00 et 0xFF.
SourceAddress	Adresse MAC du nœud
SupEtherType	EtherType 0x88FB réservé pour le protocole PRP.
SupPath	Réservé, défini sur 0
SupVersion	Version du protocole, définie sur "1" pour PRP. La mise en œuvre de la version X du protocole doit interpréter une version >X comme s'il s'agissait de la version X en ignorant tous les paramètres et/ou indicateurs ajoutés par la version plus récente et interpréter les trames PRP_Supervision de la version <=X exactement comme spécifié pour la version concernée.
SupSequenceNumber	Numéro de séquence incrémenté de 1 pour chaque trame de supervision envoyée.
TLV1.Type	Mode d'exploitation. Les valeurs 20 et 21 indiquent respectivement que le nœud prend en charge le rejet des doublons et qu'il met en œuvre l'acceptation des doublons. Les autres valeurs sont réservées.
TLV1.Length	Longueur de l'adresse MAC suivante, en octets.
MacAddress	Adresses MAC utilisées par les deux ports du nœud.
TLV2.Type	Absent si le nœud n'est pas une RedBox
TLV2.Length	Absent si le nœud n'est pas une RedBox
RedBoxMacAddress	Absent si le nœud n'est pas une RedBox
TLV0.Type	Fermeture de TLV, défini sur 0
TLV0.Length	Fermeture de TLV, défini sur 0

4.3.3 Trame PRP_Supervision pour RedBox

Une RedBox, c'est-à-dire un nœud qui fait office de proxy pour un ou plusieurs SAN (appelés VDAN ou DAN virtuels), doit ajouter au champ TLV1 un champ TLV2 supplémentaire avec le contenu donné dans le Tableau 7:

Tableau 7 – TLV PRP_Supervision pour RedBox

Nom de champ	Description
TLV1.Type	Mode d'exploitation. La valeur 20 indique que le nœud prend en charge le mode de rejet des doublons, car il s'agit du seul mode qu'une RedBox prend en charge. Les autres valeurs sont réservées.
TLV1.Length	Longueur de l'adresse MAC suivante, en octets.
MacAddress	Adresse MAC du VDAN
TLV2.Type	Identificateur de la RedBox, la valeur 30 indiquant que le nœud est une RedBox
TLV2.Length	Longueur de l'adresse MAC suivante, en octets.
RedBoxMacAddress	Adresse MAC de la RedBox qui fait office de proxy pour l'autre appareil (VDAN)
TLV0.Type	Fermeture de TLV, défini sur 0
TLV0.Length	Fermeture de TLV, défini sur 0

4.3.4 Réception d'une trame PRP_Supervision et NodesTable

Lors de la réception d'une trame PRP_Supervision sur n'importe quel LAN, un nœud doit créer une entrée dans le NodesTable correspondant au MacAddress de cette source, comme indiqué dans le corps du message et non dans l'adresse source, avec le mode d'acceptation ou de rejet des doublons, comme indiqué dans la trame.

Si un nœud cesse de recevoir des trames PRP_Supervision d'une source pendant une période supérieure à NodeForgetTime, mais reçoit des trames de cette source sur un LAN seulement, il doit remplacer le statut de ce nœud par SanA ou SanB, selon le LAN d'origine des trames.

NOTE 1 Cette règle permet de déplacer un SAN entre LAN_A et LAN_B et d'obtenir le mode d'exploitation pour un SAN s'il a d'abord été enregistré à l'envoi et non à la réception, car un DANP commence par effectuer des envois sur les deux LAN.

NOTE 2 Cette règle permet de faire la distinction entre un SAN et un DANP en mode d'acceptation des doublons lorsqu'une ligne est déconnectée.

4.4 Nœud de pontage

Si le paramètre "nœud de pontage" est activé, le nœud doit faire office de nœud de pontage pour ses deux ports, conformément à un protocole non spécifique, c'est-à-dire HSR, RSTP ou MRP, mais ne doit pas ajouter de RCT aux trames qu'il envoie.

NOTE 1 Le paramètre de nœud de pontage prend en charge l'association d'un DANP à 2 ports du LAN pour mettre en œuvre une topologie à redondance partielle. Il n'y a pas d'exigence pour ponter les trames normalisées en cas de défaillance double, mais les personnes chargées de la mise en œuvre sont libres de l'ajouter.

NOTE 2 La présence de nœuds en mode de pontage présente un danger potentiel pour un réseau PRP, puisqu'il va à l'encontre de l'hypothèse d'indépendance en cas de défaillance des deux LAN.

4.5 Constantes

Les paramètres de constantes sont donnés dans le Tableau 8.

Les autres valeurs peuvent être définies par l'utilisateur.

Tableau 8 – Constantes PRP

Constante	Description	Valeur par défaut
LifeCheckInterval	Intervalle type entre deux trames PRP_Supervision successives envoyées par le même nœud.	2 000 ms
NodeForgetTime	Temps à l'issue duquel une entrée de nœud est effacée de NodesTable après que les trames de ce nœud ont cessé d'être reçues.	60 000 ms
EntryForgetTime	Temps maximal pendant lequel une entrée peut résider dans la table des doublons, indépendamment de la vitesse.	400 ms
NodeRebootInterval	Temps minimal pendant lequel un nœud qui se réinitialise reste silencieux, indépendamment de la vitesse.	500 ms

4.6 Spécification de service PRP

On peut accéder aux objets de la LRE, les compteurs d'erreurs par exemple, à partir d'une application qui fonctionne sur le même nœud qui utilise les appels directs ou à partir d'une application présente sur un autre nœud qui utilise un protocole de gestion comme SNMP ou un protocole particulier comme l'IEC 61850.

Tous les objets LRE sont décrits à l'Article 7 comme une base d'informations de gestion (MIB) conforme à SNMP, (RFC 2578, RFC 3418), commune à PRP et HSR.

5 Redondance transparente haute disponibilité (HSR)

5.1 Objectifs du protocole HSR

L'Article 5 décrit l'application des principes du protocole PRP de l'Article 4 pour mettre en œuvre une redondance transparente haute disponibilité (HSR), qui conserve la propriété PRP de récupération instantanée, applicable à toute topologie, notamment dans les anneaux et les anneaux d'anneaux.

NOTE Par rapport au PRP, le protocole HSR permet de réduire environ de moitié l'infrastructure réseau. Pour les anneaux basés sur les normes IEEE 802.1D (RSTP), IEC 62439-2 (MRP), IEC 62439-6 (DRP) ou IEC 62439-7 (RRP), le trafic du réseau est à peu près le même.


Les nœuds de l'anneau ne sont que des nœuds de pontage compatibles HSR, ce qui permet d'éviter l'utilisation de ponts dédiés. Les nœuds à une seule association (SAN), comme les ordinateurs portables ou les imprimantes, ne peuvent pas être associés directement à l'anneau. Ils doivent être associés à un boîtier de redondance (RedBox).

5.2 Principe de fonctionnement du protocole HSR

5.2.1 Fonctionnement de base avec une topologie en anneau

Comme dans le PRP, un nœud dispose de deux ports qui fonctionnent en parallèle; il s'agit d'un DANH (Doubly Attached Node avec protocole HSR).

Un réseau HSR simple se compose de nœuds de pontage à deux connexions. Chaque nœud dispose de deux ports d'anneau, interconnectés par des liaisons full duplex, comme présenté dans l'exemple de la Figure 10 (multidiffusion) et de la Figure 11 (monodiffusion) pour une topologie en anneau.



IEC

Anglais	Français
bridge	pont
"A" frame	trame "A"
"B" frame	trame "B"
"C" frame	trame "C"
"D" frame	trame "D"
interlink	interconnexion
"A" frame (counter-clockwise)	trame "A" (sens inverse des aiguilles d'une montre)
"B" frame (clockwise)	trame "B" (sens des aiguilles d'une montre)
"C" frame (non-HSR frame from host to ring)	trame "C" (trame non HSR de l'hôte à l'anneau)
"D" frame (non-HSR frame from ring to host)	trame "D" (trame non HSR de l'anneau à l'hôte)
frame removed from the ring by the node	trame retirée de l'anneau par le nœud

Figure 10 – Exemple HSR d'une configuration en anneau pour le trafic multidiffusion


Un DANH source envoie une trame transmise depuis ses couches supérieures (trame "C"), la préfixe avec une balise HSR pour identifier les trames en double, puis envoie la trame vers chaque port (trame "A" et trame "B").

Un DANH de destination reçoit, dans un état sans panne, deux trames identiques de chaque port dans un intervalle donné, supprime la balise HSR de la première trame avant de la transmettre aux couches supérieures (trame "D") et rejette tout doublon.

Les nœuds permettent la fonctionnalité de pont IEEE 802.1D et transmettent les trames d'un port au suivant, conformément à ces quatre règles:

- 1) un nœud ne transmettra pas une trame qu'il a injectée dans l'anneau ou la maille;

- 2) un nœud ne transmettra pas une trame de laquelle il est la destination unique (sauf dans des applications particulières comme la supervision de la redondance);
- 3) un port n'enverra pas de trame qui est le doublon d'une trame déjà envoyée dans la même direction;
- 4) un port s'abstiendra (facultativement) d'envoyer une trame qui est le doublon d'une trame déjà reçue dans la direction opposée (sauf pour les trames de supervision et de temporisation).



IEC

Anglais	Français
bridge	pont
"A" frame	trame "A"
"B" frame	trame "B"
"C" frame	trame "C"
"D" frame	trame "D"
interlink	interconnexion
"A" frame (counter-clockwise)	trame "A" (sens inverse des aiguilles d'une montre)
"B" frame (clockwise)	trame "B" (sens des aiguilles d'une montre)
"C" frame (non-HSR frame from host to ring)	trame "C" (trame non HSR de l'hôte à l'anneau)
"D" frame (non-HSR frame from ring to host)	trame "D" (trame non HSR de l'anneau à l'hôte)
frame removed from the ring by the node	trame retirée de l'anneau par le nœud

Figure 11 – Exemple HSR d'une configuration en anneau pour le trafic monodiffusion

Les trames circulant dans l'anneau portent la balise HSR insérée par la source, qui contient un numéro de séquence. Le doublet {adresse MAC source, numéro de séquence} identifie de manière unique les différentes copies d'une même trame.

NOTE 1 Le décalage maximal entre deux trames d'une paire dépend de la position relative du nœud de réception et du nœud d'envoi. Dans la situation la plus défavorable dans laquelle chaque nœud de l'anneau transmet au même moment sa propre trame d'une taille égale à la taille maximale de 1 536 octets (longueur maximale prise en

charge par EtherType défini dans ISO/IEC/IEEE 8802-2:1998), chaque nœud pourrait introduire un retard de 125 µs à 100 Mbit/s. Avec 50 nœuds, le déphasage peut dépasser 6 ms.


NOTE 2 Le retard minimal avant qu'une trame retourne au nœud qui l'a injectée dépend de la taille de l'anneau. Dans la situation la plus défavorable où il n'y a pas d'autre trafic et où chaque nœud de l'anneau fonctionne en mode à la volée avec le temps de résidence de 4 µs, dans un anneau à 32 nœuds, une trame d'une taille maximale de 1 536 octets (125 µs à 100 Mbit/s) retourne au nœud d'injection au bout de 124 µs, c'est-à-dire avant que le nœud d'injection ait fini de la transmettre.

5.2.2 Structure des nœuds DANH

La Figure 12 présente une vue conceptuelle de la structure d'un DANH mise en œuvre matériellement; les mises en œuvre pratiques peuvent être différentes. Les deux ports HSR A et B et le port de l'appareil C sont connectés par la LRE, qui inclut une matrice de commutation qui permet de transférer des trames d'un port à l'autre. La matrice de commutation permet un pontage à la volée. La LRE présente aux couches supérieures la même interface qu'un émetteur-récepteur Ethernet normalisé.

Le circuit d'entrée vérifie si ce nœud est la destination de la trame et procède éventuellement à un filtrage VLAN et multidiffusion pour décharger le processeur. Le rejet des doublons est mis en œuvre dans les files d'attente de sortie.

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016 PDF



Anglais	Français
upper layers	couches supérieures
applications	applications
publisher/subscriber	fournisseur/abonné
link layer interface	interface avec la couche de liaison
transport layer	couche de transport
network layer	couche réseau
HSR tagger	baliseur HSR
link redundancy entity (LRE)	LRE
forwarding	transfert
filtering	filtrage
demux	démultiplexeur
mux	multiplexeur
duplicate discard	rejet de doublons
duplicate discard and untagger	rejet de doublons et outil de débalisage

Figure 12 – Structure HSR d'un DANH

5.2.3 Topologie

5.2.3.1 Association de nœuds à une seule association


Les nœuds à une seule association (SAN), comme les ordinateurs portables de maintenance ou les imprimantes, ne peuvent pas être insérés directement dans l'anneau, car ils ne disposent que d'un seul port et ne peuvent pas interpréter la balise HSR contenue dans les trames. Les SAN communiquent avec les appareils de l'anneau par le biais d'une RedBox qui

fait office de proxy pour les SAN qui lui sont associés, comme indiqué aux Figure 10 et Figure 11. La RedBox est détaillée en 5.2.3.8.

La connexion de nœuds non HSR aux ports de l'anneau, qui ouvre l'anneau, peut permettre la configuration. Le trafic non HSR dans l'anneau fermé est pris en charge dans un mode facultatif.

5.2.3.2 Utilisation du protocole HSR avec des LAN distincts

Les nœuds HSR peuvent être connectés de la même manière que les nœuds PRP. A cet effet, un nœud HSR peut être défini sur "pas de transfert" (mode N). A la différence du PRP, les SAN ne peuvent pas être associés directement à un tel réseau dupliqué, sauf s'ils peuvent interpréter la balise HSR.



IEC

Anglais	Français
bridge	pont
bridged local area network (ring) LAN_A	LAN ponté (anneau) LAN_A
bridged local area network (tree) LAN_B	LAN ponté (arborescence) LAN_B

Figure 13 – Exemple HSR de topologie qui utilise deux réseaux indépendants

5.2.3.3 Couplage d'anneaux par homologues

Deux anneaux HSR peuvent être connectés par des appareils à quatre ports qui disposent de capacités de transfert et nommés QuadBox, comme le montre la Figure 14. Cette méthode présente des avantages lorsque le flux de trafic dépasse les capacités d'un seul anneau. Elle ne permet toutefois pas d'améliorer les retards de transmission d'une extrémité à l'autre.


Bien qu'une QuadBox suffise pour amener le trafic sans panne au réseau, deux QuadBox permettent de prévenir les points uniques de défaillance.

Une QuadBox transfère les trames sur chaque anneau comme un nœud HSR et transmet les trames non modifiées à l'autre anneau, sauf si la trame peut être identifiée comme une trame

à ne pas transférer à l'autre anneau. A cet effet, une QuadBox est censée filtrer le trafic en fonction, par exemple, d'un filtrage multidiffusion ou d'un filtrage VLAN. Il n'y a pas d'apprentissage des adresses MAC dans une QuadBox, car l'apprentissage des adresses MAC sur des ports spécifiques d'un appareil QuadBox pourrait entraîner une interruption des communications en cas de panne de la QuadBox qui a appris une adresse et qui transfère le trafic réseau.

Avec des QuadBox réalisées en tant qu'entités physiques uniques, les deux anneaux interconnectés partagent le même domaine de redondance pour la tolérance aux pannes. Si une QuadBox tombe en panne, les deux anneaux interconnectés sont dans un état dégradé et ne peuvent pas tolérer une autre panne.

La conception de QuadBox comme une RedBox peut donc contribuer à maintenir l'indépendance de la redondance. La QuadBox est alors composée de deux appareils liés par une interconnexion. Les spécifications de la RedBox incluent donc la connexion HSR.



Anglais	Français
Ring	Anneau
interlink	interconnexion
A-frame	trame A
B-frame	trame B
(next ring)	(anneau suivant)

Figure 14 – Exemple HSR du couplage par homologues de deux anneaux

En présence de deux QuadBox sur le même anneau, deux copies de la même trame sont transférées du premier anneau jusqu'au second, chacun générant deux autres copies.


Cela n'entraîne pas la diffusion de quatre trames sur le deuxième anneau, car, lorsqu'une copie issue d'une première QuadBox atteint la deuxième QuadBox sur le même deuxième anneau, la deuxième QuadBox ne la transfère pas si elle a déjà envoyé une copie issue de son interconnexion. Inversement, si la deuxième QuadBox n'a pas encore reçu de copie de son interconnexion, elle transférera la trame, mais pas la copie qui provient par la suite de l'interconnexion.

Lorsqu'une QuadBox reçoit une trame qui est elle-même injectée dans l'anneau ou une trame que l'autre QuadBox a insérée dans l'anneau, elle la transfère à l'interconnexion et à son autre port si elle n'a pas déjà envoyé une copie. Ce doublon sera rejeté à l'autre extrémité de l'interconnexion. Ce schéma peut entraîner du trafic supplémentaire sur l'interconnexion, mais il permet de simplifier la conception de la logique.

NOTE Le décalage maximal entre deux trames d'une paire est à peu près identique au décalage maximal produit si tous les nœuds résidaient sur le même anneau.

5.2.3.4 Topologie en anneau hiérarchique

Un réseau HSR peut se composer d'anneaux connectés par des QuadBox (voir Figure 15).



IEC

Anglais	Français
Ring	anneau

Figure 15 – Exemple HSR d'anneaux connectés


Bien qu'une seule QuadBox suffise pour maintenir le trafic, deux QuadBox indépendantes sont nécessaires pour éviter un point unique de défaillance.

Des SAN sont connectés directement au DANH qui fonctionne comme une RedBox simplifiée.

5.2.3.5 Connexion d'un anneau HSR à un réseau PRP

Un HSR peut être couplé à un réseau PRP par le biais de deux RedBox, une pour chaque LAN (voir Figure 16). Dans ce cas, les RedBox sont configurées pour prendre en charge le trafic PRP sur l'interconnexion et le trafic HSR sur les ports de l'anneau.

Le numéro de séquence de la RCT PRP est réutilisé pour la balise HSR et inversement pour permettre l'identification de la trame d'un réseau à l'autre et identifier les paires et doublons sur l'anneau HSR, introduits par une injection double dans l'anneau via deux RedBox HSR.



IEC

Anglais	Français
PRP nodes	nœuds PRP
end node	nœud d'extrémité
interlink	interconnexion
"A" frame (counter-clockwise)	trame "A" (sens inverse des aiguilles d'une montre)
"B" frame (clockwise)	trame "B" (sens des aiguilles d'une montre)
"C" frame (non-HSR frame from host to ring)	trame "C" (trame non HSR de l'hôte à l'anneau)
"D" frame (non-HSR frame from ring to host)	trame "D" (trame non HSR de l'anneau à l'hôte)
frame removed from the ring by the node	trame retirée de l'anneau par le nœud

Figure 16 – Exemple HSR du couplage de deux réseaux locaux PRP redondants à un anneau

Les RedBox HSR qui permettent de connecter l'anneau à un réseau PRP fonctionnent de la même manière que les RedBox utilisées pour associer les SAN décrites en 5.2.3.1, à

l'exception du fait qu'elles sont configurées en tant que RedBox "A" ou RedBox "B" pour accepter les trames PRP sur leur interconnexion.


A la Figure 16, les RedBox A et B enverraient la même trame (respectivement A et AB et B et BA), mais si une RedBox reçoit la trame avant qu'elle n'ait pu l'envoyer elle-même, elle ne l'envoie pas.

Dans l'exemple de la Figure 16, la RedBox A ne générera pas de trame "A" au nom du LAN A si elle a précédemment reçu de l'anneau la même trame "AB". Inversement, la RedBox "B" générera une trame "AB" si elle n'a pas précédemment reçu de l'anneau une trame "A", ce qui se produit dès que la trame "A" n'est pas une trame multidiffusion.

Les trames multidiffusion ou monodiffusion sans récepteur dans l'anneau (flèches vides dans la Figure 16) sont supprimées par la RedBox qui les a insérées dans l'anneau si elles provenaient de l'extérieur de l'anneau. Les trames transportent à cet effet un identificateur LAN qui est en même temps l'identificateur RedBox.

La Figure 17 présente le même couplage pour une source située dans l'anneau.

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2016



IEC

Anglais	Français
PRP nodes	nœuds PRP
end node	nœud d'extrémité
interlink A	interconnexion A
interlink B	interconnexion B
"A" frame (counter-clockwise)	trame "A" (sens inverse des aiguilles d'une montre)
"B" frame (clockwise)	trame "B" (sens des aiguilles d'une montre)
"C" frame (non-HSR frame from host to ring)	trame "C" (trame non HSR de l'hôte à l'anneau)
"D" frame (non-HSR frame from ring to host)	trame "D" (trame non HSR de l'anneau à l'hôte)
frame removed from the ring by the node	trame retirée de l'anneau par le nœud


Figure 17 – Exemple HSR de couplage du nœud d'un anneau à des LAN PRP redondants

Pour éviter de réinjecter une trame dans le réseau PRP par l'intermédiaire de l'autre RedBox, chaque trame HSR transporte l'identificateur du réseau PRP d'origine de la trame, le PathId à 4 bits. Les RedBox sont configurées et identifiées par le PathId du réseau LAN PRP auquel elles sont associées.

5.2.3.6 Connexion de plusieurs réseaux PRP à un anneau HSR

Jusqu'à 7 réseaux PRP peuvent être connectés au même anneau HSR; chacun d'eux est identifié par un NetId de 3 bits.

Les deux RedBox qui connectent un réseau PRP à un anneau HSR sont configurées avec le NetId (1 à 7) et le LanId (A=0/B=1) qui forme le PathId à 4 bit (voir Figure 18).



IEC

Anglais	Français
PRP nodes	nœuds PRP
ring	anneau

Figure 18 – Exemple HSR de couplage d'un anneau à deux LAN PRP


Pour éviter de réinjecter des trames issues d'un réseau PRP dans un autre réseau PRP ou issues de l'autre LAN dans le même réseau PRP, une RedBox assure seulement le transfert à partir des trames de l'anneau HSR qui ne transportent pas son propre PathId.

Lors de l'insertion dans l'anneau d'une trame PRP en provenance du LAN A ou du LAN B d'un réseau PRP avec un NetId donné, une RedBox insère dans le PathId de la balise HSR son propre NetId; ainsi que le LanId, autrement dit une valeur parmi "2"/"3", "4"/"5", "6"/"7", "8"/"9", "A"/"B", "C"/"D" ou "E"/"F", selon qu'il s'agisse d'une RedBox A ou d'une RedBox B.

Inversement, lors du transfert d'une trame de l'anneau à un réseau PRP, une RedBox insère le LanId "A" ou "B" dans la RCT selon qu'il s'agisse d'une RedBox A ou B.

5.2.3.7 Connexion d'un réseau PRP à plusieurs anneaux HSR

Un réseau PRP peut être connecté à un nombre quelconque d'anneaux HSR, mais ces anneaux ne peuvent pas être connectés entre eux, ni par des QuadBox, ni par un autre réseau PRP, puisque cela créerait des boucles (voir Figure 19).




IEC

Anglais	Français
PRP nodes	nœuds PRP
end node	nœud d'extrémité
Ring	anneau

Figure 19 – Exemple HSR de couplage de trois anneaux à un LAN PRP

5.2.3.8 Topologie en maille

Le protocole HSR permet toutes sortes de maillages et fournit une redondance tant que la structure ne contient aucun point unique de défaillance. Par exemple, la Figure 20 présente une disposition en matrice des nœuds. Dans ce cas, les nœuds disposent de plus de deux ports qui fonctionnent en parallèle et qui fonctionnent comme les QuadBox. Une trame reçue de l'un des ports est transférée à tous les autres ports, à l'exception de celui qui l'a reçue; chaque port transfère la trame sauf s'il a déjà envoyé un doublon.




IEC

Figure 20 – Exemple HSR de topologie en maille

5.2.4 Structure d'une RedBox

La Figure 21 présente la structure générale d'une RedBox:



IEC

Anglais	Français
to singly attached nodes	vers les nœuds à une seule connexion (SAN)
RedBox local applications	applications locales RedBox
layer 2 access	accès à la couche 2
transport layer	couche de transport
network layer	couche réseau
bridge	pont
interlink	interconnexion
send duplicate	envoi de doublons
discard duplicate	rejet de doublons
Proxy Node Table	table de nœuds proxy
link redundancy entity (LRE)	entité de redondance de liaison (LRE)
Nodes Table	table de nœuds
HSR port	port HSR

Figure 21 – Structure HSR d'une RedBox

La RedBox dispose d'une LRE qui accomplit les fonctions du protocole HSR, en particulier:

- elle transfère les trames reçues d'un port HSR à l'autre port HSR, sauf si la trame a déjà été envoyée;
- elle reçoit les trames adressées à ses propres protocoles supérieurs;
- elle préfixe les trames envoyées par ses propres couches supérieures avec la balise HSR correspondante avant d'envoyer deux copies via ses ports HSR.

Une RedBox peut être utilisée dans l'un des trois modes suivants: en tant que connexion SAN, PRP ou HSR (moitié d'une QuadBox). Selon le mode d'exploitation, la gestion des trames au niveau de l'interface d'interconnexion de la RedBox diffère (voir 5.4.2 et 5.4.4).

La logique de commutation de la Figure 21 peut être intégrée à la RedBox, de sorte que l'interconnexion devienne une connexion interne.

Une RedBox simple est présente dans chaque nœud, car la LRE fait une transition vers un hôte unique non-HSR. De plus, il est courant de disposer de plusieurs hôtes dans un nœud, car il existe souvent un port de maintenance.

5.3 Spécifications du nœud HSR

5.3.1 Fonctionnement HSR

5.3.1.1 Propriétés du nœud

Un DANH doit pouvoir fonctionner dans l'un des modes suivants, qui doit pouvoir être changé lors de l'exécution à l'aide de commandes de gestion:

- Mode H (obligatoire, mode par défaut): transfert avec balise HSR – dans ce mode, le DANH doit insérer la balise HSR au nom de son hôte et doit transférer le trafic de l'anneau, à l'exception des trames envoyées par le nœud lui-même, des trames en doublon et des trames dont le nœud est la destination unique.
- Mode N (facultatif): pas de transfert – dans ce mode, le nœud doit se comporter comme en mode H, sauf qu'il ne doit pas transférer le trafic de l'anneau entre les ports.
- Mode T (facultatif): transfert transparent – dans ce mode, le DANH doit supprimer la balise HSR avant de transférer la trame à l'autre port et doit envoyer une trame de l'hôte aux deux ports, sans insérer de balise ni rejeter les doublons.
- Mode M (facultatif): transfert mixte avec balise HSR et non HSR – dans ce mode, le DANH doit insérer la balise HSR en fonction de critères locaux lors de l'injection des trames dans l'anneau. Les trames qui disposent d'une balise HSR et qui sont issues des ports de l'anneau sont gérées comme dans le Mode H. Les trames qui disposent d'une balise non HSR et qui sont issues des ports de l'anneau sont gérées conformément aux règles de transfert de l'IEEE 802.1D.
- Mode U (facultatif): transfert monodiffusion – dans ce mode, le nœud doit se comporter comme en Mode H, sauf qu'il doit aussi transférer le trafic dont il est la destination unique.
- Mode X (facultatif): pas d'envoi vers le doublon homologue – dans ce mode, le nœud doit se comporter comme en Mode H, mais un port ne doit pas envoyer de trame qui est un doublon d'une trame entièrement et correctement reçue de la direction opposée.

5.3.1.2 Numéro de séquence de l'hôte

Un nœud HSR doit maintenir un numéro de séquence pour chaque adresse MAC utilisée par l'hôte. Le numéro de séquence commence à 0.

NOTE La LRE est censée détecter l'adresse MAC de l'hôte en écoutant les trames que ce dernier lui envoie. L'adresse MAC de l'hôte peut aussi être configurée dans la LRE. Si elle ne connaît pas l'adresse MAC de l'hôte, la LRE transmet tout le trafic HSR à l'hôte en traitant les trames monodiffusion comme des trames multidiffusion.

5.3.1.3 Adresse multidiffusion de l'hôte

Un nœud doit pouvoir filtrer au moins des adresses multidiffusion MulticastFilterSize (voir Tableau 11) ou des groupes d'adresses multidiffusion. Si la table n'est pas configurée, la LRE doit transférer toutes les trames multidiffusion à l'hôte et à l'autre port, sauf les doublons.

5.3.1.4 Adresse monodiffusion de l'hôte

Un nœud doit pouvoir communiquer ses adresses MAC à la LRE. Si l'adresse MAC n'est pas configurée, la LRE doit transférer toutes les trames monodiffusion à l'hôte et à l'autre port, sauf les doublons.

5.3.1.5 Réinitialisation de l'hôte

Lors de sa réinitialisation, un nœud doit éviter d'envoyer des trames vers ses ports HSR pendant une durée égale à NodeRebootTime (voir Tableau 11) et préférentiellement commencer la transmission en envoyant une trame de supervision.

5.3.2 Réception de données par le DANH depuis son interface de couche de liaison

Pour que chaque trame envoie au nom de son interface de couche de liaison, un nœud source doit procéder comme suit:

If the node is in Mode T		
	do not modify the frame;	
	duplicate the frame, enqueue it for sending into both HSR ports	
Else if the node is in Mode M		
	If the local criteria is not met	
		insert the HSR tag and increment the sequence number depending on local criteria; duplicate the frame, enqueue it for sending into both HSR ports
	Else	
		do not insert the HSR tag; process the frame according to IEEE 802.1D rules
Else		
	insert the HSR tag with the sequence number of the host; increment the sequence number, wrapping through 0; duplicate the frame, enqueue it for sending into both HSR ports	

NOTE 1 La mise en file d'attente signifie que la trame sera envoyée dès lors que la file d'attente ne contient plus de trame ou de trame à priorité plus élevée et que le support est prêt.

NOTE 2 L'envoi d'une trame non-HSR aux deux ports n'entraîne pas la circulation de trames, car ces trames ne seront pas transférées par le nœud HSR adjacent. Ce mécanisme est destiné à la configuration hors de l'anneau d'un nœud HSR via un PC classique.

5.3.3 Réception de données par le DANH depuis un port HSR

Un nœud qui reçoit une trame depuis l'un de ses ports HSR doit procéder comme suit:

If this frame is not HSR-tagged:		
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)	
		consume the frame and do not forward it.
	Else if the node is in Mode M	
		forward according to 802.1D forwarding rules
	Else if the node is in Mode T	
		enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.
	Else	
		enqueue the unchanged frame for passing to its link layer interface. do not forward this frame
Else (HSR-tagged frame):		
	If the node is in Mode T	
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface
	Else (node is not in Mode T)	
	If this node is a destination:	
		If this is the first occurrence of the frame over the link layer interface:
		remove the HSR tag and pass the modified frame to its link layer interface.
		Else (this is not the first occurrence of the frame over the link layer interface):
		do not pass the frame to the link layer interface.
	Else (if this node is not a destination):	
		do not pass the frame to the link layer interface.
	If this node is not the only destination or the node is in Mode U:	
		If this is the first occurrence of the frame and the node is not in Mode N:
		enqueue the unmodified frame for sending over the other port(s)
		Else (this is not the first occurrence of the frame):
		discard the frame.
	Else (if this node is not in Mode U and it is the unique destination):	
		discard the frame.

NOTE 1 Plusieurs doublons peuvent arriver, notamment lorsque les anneaux sont couplés.

NOTE 2 Un nœud accepte aussi une trame avec une balise HSR si le LanId ne correspond pas au PortId et si le LSDUsize ne correspond pas à la taille de trame.

5.3.4 Règles de transfert du DANH

Un nœud ne doit pas envoyer vers un port une trame qui est un doublon d'une trame précédemment envoyée vers ce port dans la même direction.

Un nœud qui reçoit une trame ne doit pas modifier son adresse source ni sa balise HSR lors de la transmission de la trame même s'il a effectivement modifié le contenu de la trame.

Un nœud qui détecte, grâce à la supervision de la qualité du signal ou la FCS, que la trame est endommagée ou tronquée, ne doit pas transférer cette trame. Toutefois, si le nœud qui fonctionne à la volée a déjà commencé le transfert, puis s'il détecte que la trame est

endommagée ou tronquée, il doit lui ajouter la séquence d'erreur décrite dans l'ISO/IEC/IEEE 8802-3:2014, 27.3.1.2.2, puis arrêter la transmission de cette trame.

Si un port connecté précédemment est déconnecté du réseau, un nœud doit purger le tampon du port pour qu'il ne puisse pas envoyer une trame obsolète, puis reprendre la mise en tampon lorsque le port est de nouveau connecté.

Pour les trames qui proviennent d'un DANH, la valeur 0 ou 1 peut être attribuée à LanId pour indiquer le port d'envoi (A = 0, B = 1).

Il n'est pas nécessaire de nettoyer les files d'attente de transmission des entrées issues de ce nœud, car une file d'attente de transmission se vide d'elle-même, que la liaison associée soit opérationnelle ou non.

Si un port a reçu une trame entière et correcte, il pourrait en principe s'abstenir d'envoyer un doublon de cette trame dans la direction opposée, puisque, par hypothèse, tous les nœuds de cette direction ont déjà correctement reçu la trame, puisqu'ils l'ont transmise. Même si cette méthode économise de la bande passante, elle présente plusieurs inconvénients:

- Cette méthode ne peut être appliquée aux trames de supervision, afin de conserver la vérification d'intégrité de l'anneau dans les deux directions.
- Cette méthode ne peut être appliquée aux trames PTP (voir Annexe A), puisque le retard est différent dans les deux directions.
- Cette méthode n'est pas applicable quand le mode "U" est activé afin de permettre à tous les nœuds de recevoir tout le trafic.
- Cette méthode n'est pas efficace pour le trafic qui n'est pas transféré par une destination unique.
- Cette méthode réduit le trafic de diffusion et de multidiffusion seulement quand le nœud ne fonctionne pas en mode à la volée.
- Cette méthode introduit un retard variable de réponse selon la position du nœud dans l'anneau et les conditions de course; on ne peut pas prévoir quel nœud, s'il y en a un, supprime le doublon.
- Cette méthode ne garantit pas que la bande passante gagnée reste disponible pour un autre trafic.
- Cette méthode n'améliore pas le retard du réseau dans le cas le plus défavorable quand il y a perte de nœud ou de trame. Les nœuds qui ont ce comportement pourraient interagir avec des nœuds qui ne l'ont pas.

La capacité d'empêcher l'envoi d'un doublon homologue est exprimée en Mode_X dans le PICS.

NOTE 1 Ces règles suppriment les trames HSR en circulation et ouvrent l'anneau, comme le fait le RSTP ou un protocole similaire. Il s'applique aux trames issues à l'origine du nœud et aux trames qui circulent par suite de la suppression d'un appareil après l'envoi d'une trame et à la fermeture de l'anneau, par un appareil de pont mécanique ou lorsqu'un DANH est mis hors tension, par exemple. Dans un anneau de 50 nœuds, une trame retourne à son émetteur au bout de 6 ms.

NOTE 2 Ces conditions permettent à un nœud de fonctionner en mode enregistrement et retransmission ou en mode à la volée. Le fait de retarder le transfert d'une trame n'affecte pas le retard d'anneau le plus défavorable.

NOTE 3 La méthode de rejet des doublons du PRP n'est pas une méthode préférentielle de rejet des doublons dans HSR, car le protocole HSR a pour objectif d'empêcher les doublons de circuler.

NOTE 4 Le fait que les numéros de séquence des trames envoyées par une source n'augmentent pas de façon monotone ne constitue pas une raison pour rejeter la trame. Cette observation peut toutefois être utilisée pour la surveillance du réseau.

NOTE 5 Dans le fonctionnement en mode à la volée, la réception complète de la balise HSR prend environ entre 2,4 µs et 100 Mbit/s, période à l'issue de laquelle le nœud décide de transférer ou pas la trame. En mode

d'enregistrement et de retransmission, la réception complète d'une trame dure environ 123 µs à 100 Mbit/s pour une longueur de trame maximale (1 522 octets).

5.3.5 CoS

Les priorités et les VLAN ne sont pas exigés pour le fonctionnement du HSR.

Un nœud HSR est attendu, comme indiqué dans son PICS:

- pour prendre en charge au moins 2 niveaux de priorité, conformément à la norme IEEE 802.1D (IEEE 802.1p);
- pour filtrer le trafic du VLAN, conformément à la norme IEEE 802.1Q;
- pour filtrer le trafic multidiffusion.

5.3.6 Synchronisation de l'horloge

La gestion des horloges selon l'IEC 61588 est décrite à l'Annexe A.

5.3.7 Accès déterministe au support

L'accès déterministe au support n'est pas spécifié.

NOTE L'accès déterministe au support peut tirer avantage de la synchronisation précise de l'horloge, puisqu'il laisse tous les nœuds commencer à transmettre des trames qui ont des données temporellement critiques en même temps dans un intervalle de temps réservé et qu'il assure que toutes les autres trames ont été entièrement transmises avant le début de cet intervalle de temps. La taille d'une trame figure dans l'en-tête HSR.

5.4 Spécifications de la RedBox HSR

5.4.1 Propriétés de la RedBox

Une RedBox est un appareil qui possède au moins trois ports: deux d'entre eux sont des ports de l'anneau pour le protocole HSR et le troisième est connecté à une interconnexion.

Une RedBox doit se comporter comme un DANH pour tout le trafic dont elle est la source ou la destination.

Outre le comportement du DANH, celui de l'interconnexion de la RedBox est décrit en 5.4.2. La RedBox fait la distinction, à la réception des trames issues du réseau HSR, entre ses interfaces internes et son interconnexion. Le comportement d'une RedBox à la réception de trames issues du réseau HSR est donc décrit en 5.4.4.

Une RedBox doit comporter un NetId configurable composé de trois bits, comme défini en 5.7.1. Il est utilisé dans la balise HSR pour les trames reçues du port d'interconnexion de la RedBox. Une connexion RedBox HSR à un LAN PRP ne transférera pas les trames HSR du NetId configuré dans ce même LAN.

Le NetId de la RedBox ne doit pas être utilisé dans la balise HSR des trames générées dans la RedBox. Ce type de trame doit plutôt utiliser un NetId de valeur nulle dans la balise HSR ou un NetId d'anneau qui peut être configuré séparément, avec une valeur par défaut nulle.

Une RedBox doit être configurable pour un de ces trois modes:

- 1) HSR-SAN: le trafic sur l'interconnexion n'est ni HSR ni PRP;
- 2) HSR-PRP: le trafic sur l'interconnexion est PRP avec une étiquette "A" ou "B";
- 3) HSR-HSR: le trafic sur l'interconnexion a une étiquette HSR.

NOTE 1 Une RedBox met elle-même en œuvre au moins une interface de gestion.

NOTE 2 Une RedBox est censée disposer de sa propre adresse IP, en particulier pour les messages de configuration. On peut y accéder via l'interconnexion ou les ports HSR.

NOTE 3 L'interconnexion peut être une connexion interne si la RedBox fait aussi office de pont.

5.4.2 Réception RedBox via l'interconnexion

Quand elle reçoit une frame via son port d'interconnexion, une RedBox doit procéder comme suit:

If the frame carries a HSR tag:			
	If the node is in Mode T		
		remove the HSR tag enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else if the RedBox operates in HSR-HSR mode		
		If the RedBox is a destination of the frame	
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else (If this is the first occurrence of the frame at the link layer interface)
			remove the HSR tag enqueue to the link layer interface of the RedBox
	If the frame is to be injected into the ring (RedBox is not only destination, Multicast/VLAN is ok)		
			If this is not the first occurrence of the frame
			discard the frame (already sent over that port)
			Else (If this is the first occurrence of the frame)
			enqueue the unmodified frame into each HSR port
	Else (If the RedBox does not operate in HSR-HSR mode)		
		discard the frame	
Else if the frame carries a PRP RCT and the RedBox is in HSR-PRP mode			
	If the node is in Mode T		
		remove the PRP RCT enqueue the modified frame to the link layer interface of the RedBox enqueue the modified frame into each HSR port	
	Else		
		If the source MAC address is not already registered:	
			Create an entry in the ProxyNodeTable;
		If the PRP tag does not correspond to the LanId of the RedBox "A" or "B"	
			increment the error counter lreCntWrongLanIdC;
	If the RedBox is a destination of the frame		
			If this is not the first occurrence of the frame at the link layer interface
			discard the frame
			Else
			enqueue to the link layer interface of the RedBox (with the PRP RCT)
	If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok)		
		If this is not the first occurrence of the frame at each HSR port	
			discard the frame (already sent over that port)
		Else (If this is the first occurrence of the frame at each HSR port)	

			reuse the PRP sequence number and the RedBox's path identifier to build the HSR tag enqueue the frame into each HSR port
Else (if the frame carries neither a HSR tag nor a PRP RCT)			
	If the node is in Mode T		
			enqueue the unmodified frame to the link layer interface of the RedBox enqueue the unmodified frame into each HSR port
	Else if the node is in Mode M		
			forward according to 802.1D forwarding rules
	Else		
			If the source MAC address is not already registered in the ProxyNodeTable:
			create an entry in the ProxyNodeTable;
			Else (If the source is already registered)
			(void)
			If the RedBox is a destination of the frame:
			enqueue to the link layer interface of the RedBox
			If the frame is to be injected into the ring (RedBox is not sole destination and multicast/VLAN is ok))
			append the HSR tag using the sequence number of that node increment the sequence number of that source; enqueue the tagged frame into each HSR port

NOTE 1 La réception d'une trame HSR sur l'interconnexion est considérée comme une erreur de configuration dans le mode de connexion PRP. Si la RedBox est utilisée en tant que moitié d'une QuadBox en mode de connexion HSR, l'interconnexion ne transportera que du trafic HSR. En mode HSR-HSR, l'interconnexion est gérée comme un port HSR supplémentaire conçu pour le couplage de la QuadBox et n'utilise donc pas de ProxyNodeTable, à la différence des RedBox PRP-HSR et SAN-HSR. La RedBox met en œuvre une NodesTable pour enregistrer la présence de nœuds HSR.

NOTE 2 La RCT n'est supprimée que si le DANH identifie positivement l'émetteur en tant que DANP si le nœud comporte une NodesTable. Les nœuds dépourvus d'une NodesTable suppriment la RCT, conformément au réglage de mode.

5.4.3 Transfert RedBox sur l'anneau

En plus des règles de transfert données en 5.3.4, une RedBox ne doit pas transférer dans l'anneau une trame monodiffusion conçue pour l'un des nœuds enregistrés dans la ProxyNodeTable. Cette condition est activée par défaut et peut être désactivée en passant la RedBox en Mode U.

5.4.4 Réception RedBox via un port HSR

Une RedBox qui reçoit une trame valide sur un port HSR doit procéder comme suit:

If this frame is not HSR-tagged:			
	If the frame is a link local traffic (IEEE 802.1D Table 7-10)		
			consume the frame and do not forward it.
	Else if the node is in Mode M		
			enqueue the unchanged frame for passing to the interlink. forward according to 802.1D forwarding rules
	Else if the node is in Mode T		
			enqueue the unchanged frame for passing to its link layer interface. enqueue the unmodified frame for sending over the second port.
	Else		

		enqueue the unchanged frame for passing to its link layer interface. do not forward this frame
Else (frame is HSR-tagged):		
	If the node is in Mode T	
		remove the HSR tag send the modified frame over the second port send the modified frame to its link layer interface
	Else	
	If this node is a destination:	
		If this is the first occurrence of the frame over the link layer interface:
		remove the HSR tag and pass the modified frame to its link layer interface.
		Else (this is not the first occurrence of the frame over the link layer interface):
		do not pass the frame to the link layer interface
	Else (if this node is not a destination):	
		do not pass the frame to the link layer interface.
	If this node is not the unique destination or the node is in Mode U:	
		If this is the first occurrence of the frame and the node is not in Mode N:
		enqueue the unmodified frame for sending over the other port(s).
		Else (this is not the first occurrence of the frame or the node is in Mode N):
		do not enqueue the unchanged frame to the second HSR port.
		If this is the first occurrence of the frame in direction of the interlink:
		If the RedBox is in SAN mode:
		remove the HSR tag;
		If Source in the ProxyNodeTable
		discard the frame
		Else
		enqueue frame for passing to the interlink.
		Else if the RedBox is in PRP mode:
		If the netId matches that of the RedBox
		discard the frame
		Else
		remove the HSR tag and append the PRP RCT with LanId ("A" or "B") of the RedBox and reusing the HSR sequence number.
		enqueue frame for passing to the interlink.
		Else (if the RedBox is in HSR mode)
		do not modify the frame.
		enqueue frame for passing to the interlink.
		Else (If this is not the first occurrence of the frame in direction of the interlink):
		discard the frame
	Else (If this node is the only (unicast) destination) and the node is not in Mode U:	
		discard the frame.

NOTE Pour la connexion HSR-HSR seulement, une RedBox ne vérifie pas si la trame a été envoyée par l'un des nœuds dont elle fait office de proxy, car elle ne peut pas déterminer si la trame a pu être envoyée par une RedBox redondante.

5.4.5 Réception RedBox via son interface de couche de liaison

Une RedBox qui reçoit une trame valide de son interface de couche de liaison doit procéder comme suit:

If the node is in Mode T	
	do not modify the frame;
	duplicate the frame, enqueue it for sending into both HSR ports
Else If the node is in Mode M	
	insert the HSR tag and increment the sequence number depending on local criteria forward depending on local criteria
Else	
	If the RedBox is in SAN mode:
	enqueue frame for passing to the interlink.
	Else If the RedBox is in PRP mode:
	insert the PRP RCT with the sequence number of the host; enqueue frame for passing to the interlink.
	Else (if the RedBox is in HSR mode)
	insert the HSR tag with the sequence number of the host; enqueue frame for passing to the interlink.
	duplicate the frame, enqueue it for sending into both HSR ports

5.4.6 Gestion de ProxyNodeTable de la RedBox

Une RedBox doit disposer d'une ProxyNodeTable avec une entrée pour chaque association représentée et qui prend en charge au moins les entrées ProxyNodeTableMaxEntries.

Une RedBox doit purger l'entrée d'un nœud dans ProxyNodeTable avec un délai d'expiration configurable (où ProxyNodeTableForgetTime est la valeur par défaut) lorsqu'elle ne reçoit pas de trame de ce nœud.

NOTE La taille réelle de ProxyNodeTable est indiquée dans le PICS.

5.4.7 CoS RedBox

Comme pour 5.3.5, sauf que la RedBox est censée prendre en charge un plus grand nombre de VLAN et des options plus complexes et complètes en matière d'ingénierie du trafic dans le réseau sur les couches de protocole supérieures, les règles de filtrage multidiffusion, par exemple, qu'un simple nœud.

5.4.8 Synchronisation de l'horloge RedBox

La gestion des horloges selon l'IEC 61588 est décrite à l'Annexe A.

5.4.9 Accès au support RedBox

Comme pour 5.3.7, sauf que la RedBox doit pouvoir agréger plusieurs nœuds pour lesquels elle fait office de proxy.

5.5 Spécification de la QuadBox

Une QuadBox doit fonctionner conceptuellement comme deux RedBox adossées en mode HSR-HSR.

5.6 Méthode de rejet des doublons

Pour le rejet des doublons, une trame doit être identifiée par:

- son adresse MAC source;
- son numéro de séquence.

NOTE Il est en outre possible d'utiliser d'autres champs de la trame, comme la somme de contrôle, pour faciliter la détection des doublons.

La méthode de rejet des doublons n'est pas spécifiée.

Une méthode de rejet des doublons doit pouvoir omettre une entrée identifiée par <adresse MAC source><numéro de séquence> après une durée EntryForgetTime.


5.7 Format de trame pour HSR

5.7.1 Format de trame pour toutes les trames

Les trames HSR sont identifiées de manière unique par leur balise HSR.

La balise HSR est composée des champs suivants:

- EtherType de 16 bits (HSR_EtherType = 0x892F)
- Identificateur de chemin de 4 bits (PathId)
- Taille de trame de 12 bits (LSDUsize)
- Numéro de séquence de 16 bits (SeqNr)



IEC

Anglais	Français
octet position	position d'octet
preamble	préambule
LSDU size	taille LSDU
payload	données utiles
HSR tag	balise HSR
original LPDU	LPDU d'origine
time	temps

Figure 22 – Trame HSR sans balise VLAN

Le champ PathId de 4 bits empêche le rejet de trames issues d'un réseau PRP vers un autre réseau PRP.

Les 3 bits de poids fort du champ de chemin de 4 bits sont le NetId avec des valeurs 0 pour les trames HSR normales, 1 à 6 pour les trames issues d'un réseau PRP (6 réseaux pris en charge au maximum, susceptible d'être aussi utilisés ultérieurement pour d'autres réseaux HSR), 7 réservés.

Le bit de poids faible du champ de chemin de 4 bits est le LanId. Ce bit indique une trame issue d'un réseau PRP s'il a été reçu sur le port A ou le port B d'une RedBox.


Pour les trames issues de l'anneau HSR, la valeur 0 ou 1 peut être attribuée à LanId pour indiquer le port d'envoi (A = 0, B = 1).

L'allocation du champ PathId de 4 bits est la suivante:

- 0000 à 0001 trame issue d'un nœud HSR (A/B);
- 0010 à 1111 trames issues de 7 réseaux PRP (A/B).

Si la trame transporte une balise VLAN conforme à l'IEEE 802.1Q, la balise HSR doit être insérée après la balise VLAN.

La même trame avec une balise VLAN est présentée à la Figure 23.



IEC

Anglais	Français
octet position	position d'octet
preamble	préambule
LSDU size	taille LSDU
payload	données utiles
time	temps
VLAN tag	balise VLAN
HSR tag	balise HSR
original LPDU	LPDU d'origine

Figure 23 – Trame HSR avec balise VLAN

La taille de trame minimale (y compris l'adresse MAC de destination et la FCS) doit être de 70 octets sans balise VLAN et de 74 octets avec balise VLAN. Le remplissage doit être inséré entre les données utiles et la FCS.

NOTE 1 LSDUsize comporte la même définition que dans PRP et contient un remplissage possible.

Tableau 10 – Trame HSR_Supervision avec balise VLAN facultative

	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0						msb	U/L	I/G								
2																
4									lsb							
6						msb	U/L	0								
8																
10									lsb							
12																
14																
16																
18																
20																
22																
24																
26																
28																
30						msb	U/L	0								
32																
34																
36																
38						msb	U/L	0								
40																
42									lsb							
44																
46																
68																
70																
72																
74																

NOTE 2 Le format de la trame est presque identique à celui des trames de supervision PRP.

NOTE 3 Les trames de supervision permettent à NodesTable de conserver une trace de la présence d'autres nœuds identifiés par leur adresse MAC. Lorsqu'un nœud n'envoie pas de trame de supervision pendant un délai d'expiration donné (de l'ordre d'une minute), son adresse est supprimée de NodesTables s'il est utilisé.

5.7.2.2 Contenu de la trame HSR_Supervision

Le contenu de la trame HSR_Supervision doit être défini comme suit:

HSR_DestinationAddress	Adresse multidiffusion réservée 01-15-4E-00-01-XX. Par défaut, XX est "00", mais, en cas de conflits, XX peut être configuré pour prendre n'importe quelle valeur comprise entre 0x00 et 0xFF.
SourceAddress	Adresse MAC du nœud.
Balise VLAN (facultative)	Balise IEEE 802.1Q en cas d'utilisation.
HSR_EtherType	EtherType du protocole HSR; voir 5.7.1
PathId	Chemin sur lequel cette trame a été envoyée, en particulier, identification du réseau PRP.
LSDUsize	Taille de la LSDU, en octets.
SeqNr	SendSeq de l'émetteur.
SupEtherType	EtherType des trames de supervision, doit être 0x88FB
SupPath	Non utilisé, défini sur 0
SupVersion	Version du protocole, définie sur "1" (un) pour cette version de HSR. La mise en œuvre de la version X du protocole doit interpréter une version >X comme s'il s'agissait de la version X en ignorant tous les paramètres et/ou indicateurs ajoutés par la version plus récente et interpréter la version <=X exactement comme spécifié pour la version concernée. La version ne doit pas dépasser 64.
SupSequenceNumber	Numéro de séquence des trames de supervision
TLV1.Type	Prend la valeur 23 pour indiquer qu'il s'agit d'un nœud HSR en fonctionnement normal. Les autres valeurs sont réservées.
TLV1.Length	Indique la longueur de l'adresse MAC suivante.
MacAddress	Adresse MAC utilisée par l'émetteur.
TLV2.Type	Identificateur du nœud source défini sur 30 si le nœud source est une RedBox (voir 4.3.3)
TLV2.Length	Indique la longueur de l'adresse MAC suivante.
RedBoxMacAddress	Adresses MAC utilisées par la RedBox sur son interconnexion.
TLV0.Type	Fermeture de TLV, défini sur 0
TLV0.Length	Fermeture de TLV, défini sur 0

5.7.2.3 Réception d'une trame HSR_Supervision et NodesTable

Lors de la réception de la première trame HSR_Supervision sur un port de l'anneau, un nœud doit créer une entrée dans la NodesTable correspondant à la MacAddress de cette source, indiquée dans le corps du message et non dans l'adresse source de la trame, puis enregistrer le port d'origine de cette trame.

La réception des trames suivantes doit être enregistrée pour identifier les erreurs de réception.

5.7.2.4 Non-réception d'une trame HSR_Supervision

Lorsqu'un appareil détecte qu'un autre appareil n'existe pas, il doit, si cet appareil reprend la transmission, accepter ses trames comme s'il n'avait pas été connecté auparavant.

NOTE Le mécanisme de détection des appareils inexistants n'est pas spécifié, il peut être pris comme hypothèse que la lecture de NodesTable réinitialise ou lise un compteur de trames.

5.8 Constantes

Les paramètres de constantes sont donnés dans le Tableau 11.

Les autres valeurs peuvent être définies par l'utilisateur.

Tableau 11 – Constantes HSR


Constante	Description	Valeur par défaut
LifeCheckInterval	Fréquence d'envoi par un nœud d'une trame HSR_Supervision	2 000 ms
NodeForgetTime	Durée à l'issue de laquelle une entrée de nœud est effacée de NodesTable	60 000 ms
ProxyNodeTableForgetTime	Durée à l'issue de laquelle une entrée de nœud est effacée de ProxyNodeTable	60 000 ms
ProxyNodeTableMaxEntries	Nombre maximal d'entrées dans ProxyNodeTable	512
EntryForgetTime	Durée à l'issue de laquelle une entrée est supprimée des doublons	400 ms
NodeRebootInterval	Durée minimale pendant laquelle un nœud qui redémarre reste silencieux	500 ms
MulticastFilterSize	Nombre d'adresses multidiffusion à filtrer	16

5.9 Spécification de service HSR

On peut accéder aux objets de la LRE, les compteurs d'erreurs par exemple, à partir d'une application qui fonctionne sur le même nœud qui utilise les appels directs ou à partir d'une application présente sur un autre nœud qui utilise un protocole de gestion comme SNMP ou un protocole particulier comme l'IEC 61850.

Tous les objets LRE sont décrits à l'Article 7 comme une MIB conforme à SNMP, (RFC 2578, RFC 3418), commune à PRP et HSR.


Pour mieux appréhender les compteurs MIB, la Figure 24 représente la structure d'un nœud avec les entités qui restent dans le compteur.



Anglais	Français
publisher/subscriber	fournisseur/abonné
transport layer	couche de transport
network layer	couche réseau
untagger	outil de débalisage
link redundancy entity (LRE)	LRE
forwarding	transfert
demux	démultiplexeur
mux	multiplexeur
dupl. discard	rejet de doublons
HSR tagger	baliseur HSR
Nodes Table	table de nœuds

Figure 24 – Nœud HSR avec compteurs de gestion

La Figure 25 présente la base d'informations de gestion d'une RedBox.



Anglais	Français
to singly attached nodes	vers les nœuds à une connexion (SAN)
RedBox local applications	applications locales RedBox
layer 2 access	accès à la couche 2
transport layer	couche de transport
network layer	couche réseau
bridge	pont
send duplicate	envoi de doublons
discard duplicate	rejet de doublons
Proxy Node Table	table de nœuds proxy
link redundancy entity (LRE)	LRE
Nodes Table	table de nœuds
HSR port	port HSR

Figure 25 – RedBox HSR avec compteurs de gestion

6 Enoncé de conformité de mise en œuvre de protocole (PICS)

Le PICS doit indiquer si les options suivantes sont prises en charge:

- SNMP_MIB: capacité à prendre en charge la MIB du SNMP
- NTAB_SIZ: nombre d'entrées dans la NodesTable (0 = pas de NodesTable)
- PRIO_QTY: nombre de priorités prises en charge
- VLAN_QTY: nombre de VLAN pris en charge
- MULT_QTY: nombre d'adresses multidiffusion filtrées
- CLK_1588: prise en charge de la synchronisation IEC 61588

- PRP_SRP: capacité d'exécution en tant que nœud sans effet de pont sans PRP
- PRP_RSTP: capacité à fonctionner en tant qu'élément de pont RSTP avec un rôle de port désigné
- PRP_MRP: capacité à fonctionner comme élément de pont MRP (client ou maître)
- HSR_H: capacité à prendre en charge le Mode H de la HSR (par défaut)
- HSR_N: capacité à prendre en charge le Mode N de la HSR (pas de transfert)
- HSR_M: capacité à prendre en charge le Mode M de la HSR (transfert mixte)
- HSR_T: capacité à prendre en charge le Mode T de la HSR (transfert transparent)
- HSR_U: capacité à prendre en charge le Mode U de la HSR (le transfert a aussi lieu en cas de destination unique)
- HSR_X: capacité à prendre en charge le Mode X de la HSR (pas de transfert vers un doublon homologue)
- RBX_PRP: RedBox avec ports PRP
- RBX_HSR: RedBox avec ports HSR
- QBX_HSR: QuadBox qui intègre deux RedBox
- RBX_PNT: nombre d'entrées dans la ProxyNodeTable

7 Base d'informations de gestion (MIB) PRP/HSR

Les objets MIB reflètent les arguments des paramètres de service qui portent le même nom et commencent par une majuscule. Si la valeur true est attribuée à l'option PICS SNMP_MIB, les structures de données MIB définies à l'Article 7 doivent être disponibles pour OID = 1.0.62439, en plus des MIB que les ports peuvent fournir, avec la définition suivante.

NOTE Lors de la copie de la MIB à partir d'un document PDF, tous les en-têtes et pieds de page sont supprimés, puis les caractères sont convertis au format exigé par le protocole SNMP. Bien que des précautions aient été prises afin de n'utiliser aucun caractère non valide, certains peuvent apparaître dans le document du fait du processus d'édition.

```
-- *****
IEC-62439-3-MIB DEFINITIONS ::= BEGIN
-- *****
-- Imports
-- *****
IMPORTS

    MODULE-IDENTITY, OBJECT-TYPE,
    Counter32, TimeTicks, Integer32, Unsigned32          FROM SNMPv2-SMI
    OBJECT-GROUP,
    MODULE-COMPLIANCE                                   FROM SNMPv2-CONF
    TruthValue, RowStatus, MacAddress, DisplayString,
    TEXTUAL-CONVENTION                                  FROM SNMPv2-TC;

-- *****
-- Root OID
-- *****

iec62439 MODULE-IDENTITY

    LAST-UPDATED "201405220000Z" -- 2014, May 22
    ORGANIZATION "IEC/SC 65C"
    CONTACT-INFO "

        International Electrotechnical Commission
        IEC Central Office
        3, rue de Varembe
        P.O. Box 131
        CH - 1211 GENEVA 20
        Switzerland
```

```

        Phone: +41 22 919 02 11
        Fax: +41 22 919 03 00
        email: info@iec.ch
    "
DESCRIPTION    "
    This MIB module defines the Network Management interfaces
    for the redundancy protocols defined by the IEC 62439 suite.
    This MIB exposes the IEC62439-3 objects (PRP + HSR)
    "

REVISION       "201405220000Z" -- 2014, May 22
DESCRIPTION    "
    added reference to ptp MIB and lreDupListResideMaxTime and type
SecondFraction
    "

REVISION       "201202170000Z" -- February 17, 2012
DESCRIPTION    "
    Consistency brought into line to mrp, crp, brp MIBs
    "

REVISION       "201108260000Z" -- August 26, 2011
DESCRIPTION    "
    This MIB is aligned with the changes to PRP and HSR as defined in the
    Amendment to IEC 62439-3
    "

REVISION       "200811100000Z" -- November 10, 2008
DESCRIPTION    "
    Separation of IEC 62439 into a suite of documents.
    This MIB applies to IEC 62439-3, added HSR functionality
    "

REVISION       "200612160000Z" -- December 16, 2006
DESCRIPTION    "
    Initial version of the Network Management interface for the
    Parallel Redundancy Protocol
    "

::= {iso std(0) 62439 }

-- *****
-- Redundancy Protocols
-- *****

mrp OBJECT IDENTIFIER ::= { iec62439 1 }
prp OBJECT IDENTIFIER ::= { iec62439 2 }
crp OBJECT IDENTIFIER ::= { iec62439 3 }
brp OBJECT IDENTIFIER ::= { iec62439 4 }
drp OBJECT IDENTIFIER ::= { iec62439 5 }
rrp OBJECT IDENTIFIER ::= { iec62439 6 }
ptp OBJECT IDENTIFIER ::= { iec62439 7 }

-- *****
-- Textual conventions
-- *****

SecondFraction ::= TEXTUAL-CONVENTION
    DISPLAY-HINT "d"
    STATUS current
    DESCRIPTION
        "time interval expressed in multiple of 2**16 = 15,7 microseconds
        this corresponds to the fraction of seconds in the NTP time representation
        the minimum time is 0 microseconds, the maximum 18 hours
        "
    REFERENCE "IEC 62439-3"
    SYNTAX Integer32

-- *****
-- Objects of the PRP Network Management
-- *****

```

```

linkRedundancyEntityNotifications OBJECT IDENTIFIER ::= { prp 20 }
linkRedundancyEntityObjects OBJECT IDENTIFIER ::= { prp 21 }
linkRedundancyEntityConformance OBJECT IDENTIFIER ::= { prp 22 }

-- *****
lreConfiguration OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 0 }
lreStatistics OBJECT IDENTIFIER ::= { linkRedundancyEntityObjects 1 }
lreConfigurationGeneralGroup OBJECT IDENTIFIER ::= { lreConfiguration 0 }
lreConfigurationInterfaceGroup OBJECT IDENTIFIER ::= { lreConfiguration 1 }
lreStatisticsInterfaceGroup OBJECT IDENTIFIER ::= { lreStatistics 1 }

-- *****
-- Objects for lreConfigurationGeneralGroup
-- *****

lreManufacturerName OBJECT-TYPE
    SYNTAX DisplayString
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "specifies the name of the LRE device manufacturer"
    ::= { lreConfigurationGeneralGroup 1 }

lreInterfaceCount OBJECT-TYPE
    SYNTAX Integer32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "total number of LREs present in this system."
    ::= { lreConfigurationGeneralGroup 2 }

-- *****
-- Objects for lreConfigurationInterfacesGroup
-- *****

lreConfigurationInterfaces OBJECT IDENTIFIER
    ::= { lreConfigurationInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesConfigTable***
-- *****

lreInterfaceConfigTable OBJECT-TYPE
    SYNTAX SEQUENCE OF LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "
        list of PRP/HSR LREs. Each entry corresponds
        to one PRP/HSR Link Redundancy Entity (LRE), each representing a
        pair of LAN ports A and B. Basic devices supporting PRP/HSR may
        have only one LRE and thus one entry in the table, while more complex
        devices may have several entries for multiple LREs.
    "
    ::= { lreConfigurationInterfaces 1 }

lreInterfaceConfigEntry OBJECT-TYPE
    SYNTAX LREInterfaceConfigEntry
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "
        each entry contains management information applicable to a
        particular LRE.
    "
    INDEX { lreInterfaceConfigIndex }
    ::= { lreInterfaceConfigTable 1 }

LREInterfaceConfigEntry ::=
    SEQUENCE {
        lreInterfaceConfigIndex Unsigned32,

```

```

lreRowStatus RowStatus,
lreNodeType INTEGER,
lreNodeName DisplayString,
lreVersionName OCTET STRING,
lreMacAddress MacAddress,
lrePortAdminStateA INTEGER,
lrePortAdminStateB INTEGER,
lreLinkStatusA INTEGER,
lreLinkStatusB INTEGER,
lreDuplicateDiscard INTEGER,
lreTransparentReception INTEGER,
lreHsrLREMode INTEGER,
lreSwitchingEndNode INTEGER,
lreRedBoxIdentity INTEGER,
lreEvaluateSupervision TruthValue,
lreNodesTableClear INTEGER,
lreProxyNodeTableClear INTEGER,
lreDupListResideMaxTime SecondFraction
}

lreInterfaceConfigIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each LRE."
 ::= { lreInterfaceConfigEntry 1 }

lreRowStatus OBJECT-TYPE
    SYNTAX RowStatus
    MAX-ACCESS read-create
    STATUS current
    DESCRIPTION "indicates the status of the LRE table entry"
 ::= { lreInterfaceConfigEntry 2 }

lreNodeType OBJECT-TYPE
    SYNTAX          INTEGER
                   {
prpmodel (1),
hsr (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION "specifies the operation mode of the LRE:

PRP mode 1 (1)
HSR mode (2)

Note: PRP mode 0 is considered deprecated and is not supported by this
revision of the MIB
"
 ::= { lreInterfaceConfigEntry 3 }

lreNodeName OBJECT-TYPE
    SYNTAX          DisplayString
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION "specifies this LRE's node name"
 ::= { lreInterfaceConfigEntry 4 }

lreVersionName OBJECT-TYPE
    SYNTAX OCTET STRING (SIZE(1..32))
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION "specifies the version of this LRE's software"
 ::= { lreInterfaceConfigEntry 5 }

```



```

lreMacAddress OBJECT-TYPE
    SYNTAX MacAddress
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Specifies the MAC address to be used by this LRE. MAC
addresses are
        identical for all ports of a single LRE
        "
 ::= { lreInterfaceConfigEntry 6 }

lrePortAdminStateA OBJECT-TYPE
    SYNTAX INTEGER
    {
        notActive (1),
        active (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Specifies whether the port A shall be active or not Active
through
        administrative action (Default: active).
        "
 ::= { lreInterfaceConfigEntry 7 }

lrePortAdminStateB OBJECT-TYPE
    SYNTAX INTEGER
    {
        notActive (1),
        active (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Specifies whether the port B shall be active or not Active
through
        administrative action (Default: active).
        "
 ::= { lreInterfaceConfigEntry 8 }

lreLinkStatusA OBJECT-TYPE
    SYNTAX INTEGER
    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port A"
 ::= { lreInterfaceConfigEntry 9 }

lreLinkStatusB OBJECT-TYPE
    SYNTAX INTEGER
    {
        up (1),
        down (2)
    }
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "shows the actual link status of the LRE's port B"
 ::= { lreInterfaceConfigEntry 10 }

lreDuplicateDiscard OBJECT-TYPE
    SYNTAX INTEGER
    {

```

```

doNotDiscard (1),
discard (2)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION
    "
    specifies whether a duplicate discard algorithm is used at
reception
    (Default: discard).
    "
::= { lreInterfaceConfigEntry 11 }

lreTransparentReception OBJECT-TYPE
    SYNTAX          INTEGER
                    {
removeRCT (1),
passRCT (2)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        if removeRCT is configured, the RCT is removed when
forwarding to the upper
        layers, only applicable for PRP LRE (Default: removeRCT).
        "
::= { lreInterfaceConfigEntry 12 }

lreHsrLREMode OBJECT-TYPE
    SYNTAX          INTEGER
                    {
modeh (1),
moden (2),
modet (3),
modeu (4),
modem (5)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        This enumeration is only applicable if the LRE is an HSR
bridging node or RedBox.
        It shows the mode of the HSR LRE:
        (1) Default mode: The HSR LRE is in mode h and bridges tagged HSR traffic
        (2) Optional mode: The HSR LRE is in mode n and bridging between its HSR ports
        Is disabled.
        Traffic is HSR tagged.
        (3) Optional mode: The HSR LRE is in mode t and bridges non-tagged HSR traffic
        between its HSR ports
        (4) Optional mode: The HSR LRE is in mode u and behaves like in mode h, except it
        does not remove unicast messages
        (5) Optional mode: The HSR LRE is configured in mixed mode. HSR frames are handled
        according to mode h. Non-HSR frames are handled according to
802.1D bridging rules.
        "
::= { lreInterfaceConfigEntry 13}

lreSwitchingEndNode OBJECT-TYPE
    SYNTAX          INTEGER
                    {
nonbridgingnode(1),
bridgingunspecified(2),
prpnode(3),
hsrredboxsan(4),
hsrnode(5),
hsrredboxhsr(6),
hsrredboxprpa(7),
hsrredboxprpb(8)

```

```

    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        This enumeration shows which feature is enabled in this
particular LRE:
    (1): an unspecified non-bridging node, e.g. SRP.
    (2): an unspecified bridging node, e.g. RSTP.
    (3): a PRP node/RedBox.
    (4): an HSR RedBox with regular Ethernet traffic on its interlink.
    (5): an HSR switching node.
    (6): an HSR RedBox with HSR tagged traffic on its interlink.
    (7): an HSR RedBox with PRP traffic for LAN A on its interlink.
    (8): an HSR RedBox with PRP traffic for LAN B on its interlink.
    "
::= { lreInterfaceConfigEntry 14 }

lreRedBoxIdentity OBJECT-TYPE
    SYNTAX          INTEGER
                    {
    id1a (2),
    id1b (3),
    id2a (4),
    id2b (5),
    id3a (6),
    id3b (7),
    id4a (8),
    id4b (9),
    id5a (10),
    id5b (11),
    id6a (12),
    id6b (13),
    id7a (14),
    id7b (15)
    }
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        Applicable to RedBox HSR-PRP A and RedBox HSR-PRP B. One ID
is used by
    one pair of RedBoxes (one configured to A and one configured to B)
    coupling an HSR ring to a PRP network. The integer value states the value
    of the path field a RedBox inserts into each frame it receives from its
    interlink and injects into the HSR ring. When interpreted as binary values,
    the LSB denotes the configuration of the RedBox (A or B), and the following
    3 bits denote the identifier of a RedBox pair.
    "
::= { lreInterfaceConfigEntry 15 }

lreEvaluateSupervision OBJECT-TYPE
    SYNTAX TruthValue
    MAX-ACCESS read-write
    STATUS current
    DESCRIPTION
        "
        True if the LRE evaluates received supervision frames. False
if it
    drops the supervision frames without evaluating. Note: LREs are required
    to send supervision frames, but reception is optional. Default value is dependent
    on implementation.
    "
::= { lreInterfaceConfigEntry 16 }

lreNodesTableClear OBJECT-TYPE
    SYNTAX          INTEGER
                    {
    noOp (0),
    clearNodeTable (1)
    }

```

```

MAX-ACCESS read-write
STATUS current
DESCRIPTION          "specifies that the Node Table is to be cleared"
::= { lreInterfaceConfigEntry 17}

lreProxyNodeTableClear OBJECT-TYPE
SYNTAX                INTEGER
                        {
noOp (0),
clearProxyNodeTable (1)
}
MAX-ACCESS read-write
STATUS current
DESCRIPTION          "specifies that the Proxy Node Table is to be cleared"
::= { lreInterfaceConfigEntry 18}

lreDupListResideMaxTime OBJECT-TYPE
SYNTAX                SecondFraction
UNITS                "binaryFractionOfSecond"
MAX-ACCESS read-write
STATUS current

DESCRIPTION          "the longest time an entry may reside in the duplicates list,
expressed as the number of seconds multiplied by 65536;
the default value is 26214 x 15 us, or 400 ms; too low a value can cause
broadcast storms"
DEFVAL                {26214}
::= { lreInterfaceConfigEntry 19}

-- *****
-- *** End lreInterfaceConfigTable ***
-- *****

-- *****
-- Objects for lreStatisticsInterfacesGroup
-- *****

lreStatisticsInterfaces OBJECT IDENTIFIER
::= { lreStatisticsInterfaceGroup 0 }

-- *****
-- ***Begin LRE InterfacesStatsTable ***
-- *****

lreInterfaceStatsTable OBJECT-TYPE
SYNTAX SEQUENCE OF LREInterfaceStatsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION          "
                        list of PRP/HSR LREs. Each entry corresponds
                        to one PRP/HSR Link Redundancy Entity (LRE), each representing a
                        pair of LAN ports A and B and a port C towards the application/interlink.
                        Basic devices supporting PRP/HSR may have only one LRE and thus one entry in
                        the table, while more complex devices may have several
                        entries for multiple LREs.
                        "
::= { lreStatisticsInterfaces 1 }

lreInterfaceStatsEntry OBJECT-TYPE
SYNTAX LREInterfaceStatsEntry
MAX-ACCESS not-accessible
STATUS current
DESCRIPTION          "
                        An entry containing management information applicable to a
                        particular LRE.
                        "

```

```

INDEX { lreInterfaceStatsIndex }
::= { lreInterfaceStatsTable 1 }

```

```

LREInterfaceStatsEntry ::=
SEQUENCE {
    lreInterfaceStatsIndex Unsigned32,
    lreCntTxA Counter32,
    lreCntTxB Counter32,
    lreCntTxC Counter32,
    lreCntErrWrongLanA Counter32,
    lreCntErrWrongLanB Counter32,
    lreCntErrWrongLanC Counter32,
    lreCntRxA Counter32,
    lreCntRxB Counter32,
    lreCntRxC Counter32,
    lreCntErrorsA Counter32,
    lreCntErrorsB Counter32,
    lreCntErrorsC Counter32,
    lreCntNodes Integer32,
    lreCntProxyNodes Integer32,
    lreCntUniqueA Counter32,
    lreCntUniqueB Counter32,
    lreCntUniqueC Counter32,
    lreCntDuplicateA Counter32,
    lreCntDuplicateB Counter32,
    lreCntDuplicateC Counter32,
    lreCntMultiA Counter32,
    lreCntMultiB Counter32,
    lreCntMultiC Counter32,
    lreCntOwnRxA Counter32,
    lreCntOwnRxB Counter32
}

```

```

lreInterfaceStatsIndex OBJECT-TYPE
    SYNTAX Unsigned32
    MAX-ACCESS not-accessible
    STATUS current
    DESCRIPTION "A unique value for each LRE."
    ::= { lreInterfaceStatsEntry 1 }

```

```

lreCntTxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port A that are HSR tagged or
        fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
    ::= { lreInterfaceStatsEntry 2 }

```

```

lreCntTxB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent over port B that are HSR tagged or
        fitted with a PRP
        Redundancy Control Trailer.
        Only frames that are HSR tagged or do have a PRP RCT are counted.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
    ::= { lreInterfaceStatsEntry 3 }

```

```

lreCntTxC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames sent towards the application interface of
the DANP or DANH
        or over the interlink of the RedBox. Frames with and without PRP RCT or HSR
        tag are counted, but not link-local frames.
        A frame aborted during the transmission is not counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 4 }

lreCntErrWrongLanA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames with the wrong LAN identifier received on
LRE port A.
        Initial value = 0. Only applicable to PRP ports.
        "
 ::= { lreInterfaceStatsEntry 5 }

lreCntErrWrongLanB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames with the wrong LAN identifier received on
LRE port B
        Initial value = 0. Only applicable to PRP ports.
        "
 ::= { lreInterfaceStatsEntry 6 }

lreCntErrWrongLanC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames with the wrong LAN identifier received on
the interlink of
        a RedBox. Only applicable to HSR RedBoxes in HSR-PRP configuration
        (hsrredboxprpa and hsrredboxprpb).
        "
 ::= { lreInterfaceStatsEntry 7 }

lreCntRxA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
        number of frames received on a LRE port A. Only frames that
are HSR tagged
        or fitted with a PRP Redundancy Control Trailer are counted. Frames that are
        not forwarded anywhere (e.g. because the sender of the frame is in the proxy
        node table) are counted, too. Only frames received completely and without
        error are counted.
        Initial value = 0.
        "
 ::= { lreInterfaceStatsEntry 8 }

lreCntRxB OBJECT-TYPE

```

```

SYNTAX Counter32
MAX-ACCESS read-only
STATUS current
DESCRIPTION
    "
        number of frames received on a LRE port B. Only frames that
are HSR tagged
    or fitted with a PRP Redundancy Control Trailer are counted. Frames that are
    not forwarded anywhere (e.g. because the sender of the frame is in the proxy
    node table) are counted, too. Only frames received completely and without
    error are counted.
        Initial value = 0.
    "
::= { lreInterfaceStatsEntry 9 }

lreCntRxC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
            number of frames received from the application interface of a
DANP or DANH or
            the number of number of frames received on the interlink of a RedBox. Frames
            with and without PRP RCT or HSR tag are counted, but not link-local frames.
            Only frames received completely and without error are counted.
                Initial value = 0.
            "
::= { lreInterfaceStatsEntry 10 }

lreCntErrorsA OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
            number of frames with errors received on this LRE port A.
            Initial value = 0.
        "
::= { lreInterfaceStatsEntry 11 }

lreCntErrorsB OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
            number of frames with errors received on this LRE port B.
            Initial value = 0.
        "
::= { lreInterfaceStatsEntry 12 }

lreCntErrorsC OBJECT-TYPE
    SYNTAX Counter32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "
            number of frames with errors received on the application
interface of a DANP or
            DANH or on the interlink of a RedBox.
                Initial value = 0.
            "
::= { lreInterfaceStatsEntry 13 }

lreCntNodes OBJECT-TYPE
    SYNTAX Integer32
    MAX-ACCESS read-only
    STATUS current
    DESCRIPTION
        "number of nodes in the Nodes Table."

```