



BSI Standards Publication

Industrial communication networks - High availability automation networks

Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR) (IEC 62439-3:2016)

National foreword

This British Standard is the UK implementation of EN IEC 62439-3:2018. It is identical to IEC 62439-3:2016. It supersedes BS EN 62439-3:2012, which is withdrawn.

The UK participation in its preparation was entrusted to Technical Committee GEL/65/3, Industrial communications: process measurement and control, including fieldbus.

A list of organizations represented on this committee can be obtained on request to its secretary.

This publication does not purport to include all the necessary provisions of a contract. Users are responsible for its correct application.

© The British Standards Institution 2018
Published by BSI Standards Limited 2018

ISBN 978 0 580 93197 0

ICS 35.040.01; 25.040.01; 35.100.05; 25.040.40; 35.110

Compliance with a British Standard cannot confer immunity from legal obligations.

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 31 March 2018.

Amendments/corrigenda issued since publication

Date	Text affected
------	---------------

English Version

**Industrial communication networks - High availability automation
networks - Part 3: Parallel Redundancy Protocol (PRP) and
High-availability Seamless Redundancy (HSR)
(IEC 62439-3:2016)**

Réseaux industriels de communication - Réseaux
d'automatisme à haute disponibilité - Partie 3 : Protocole de
redondance parallèle (PRP) et redondance transparente de
haute disponibilité (HSR)
(IEC 62439-3:2016)

Industrielle Kommunikationsnetze - Hochverfügbare
Automatisierungsnetze - Teil 3: Parallelredundanz-Protokoll
(PRP) und nahtloser Hochverfügbarkeits-Ring (HSR)
(IEC 62439-3:2016)

This European Standard was approved by CENELEC on 2016-05-05. CENELEC members are bound to comply with the CEN/CENELEC Internal Regulations which stipulate the conditions for giving this European Standard the status of a national standard without any alteration.

Up-to-date lists and bibliographical references concerning such national standards may be obtained on application to the CEN-CENELEC Management Centre or to any CENELEC member.

This European Standard exists in three official versions (English, French, German). A version in any other language made by translation under the responsibility of a CENELEC member into its own language and notified to the CEN-CENELEC Management Centre has the same status as the official versions.

CENELEC members are the national electrotechnical committees of Austria, Belgium, Bulgaria, Croatia, Cyprus, the Czech Republic, Denmark, Estonia, Finland, Former Yugoslav Republic of Macedonia, France, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Switzerland, Turkey and the United Kingdom.



European Committee for Electrotechnical Standardization
Comité Européen de Normalisation Electrotechnique
Europäisches Komitee für Elektrotechnische Normung

CEN-CENELEC Management Centre: Rue de la Science 23, B-1040 Brussels

European foreword

The text of document 65C/834/FDIS, future edition 3 of IEC 62439-3, prepared by IEC Subcommittee 65C "Industrial networks", of IEC/TC 65 "Industrial-process measurement, control and automation" was submitted to the IEC-CENELEC parallel vote and approved by CENELEC as EN IEC 62439-3:2018.

The following dates are fixed:

- latest date by which the document has to be (dop) 2018-08-02
implemented at national level by
publication of an identical national
standard or by endorsement
- latest date by which the national (dow) 2021-02-02
standards conflicting with the
document have to be withdrawn

This document supersedes EN 62439-3:2012.

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

Endorsement notice

The text of the International Standard IEC 62439-3:2016 was approved by CENELEC as a European Standard without any modification.

In the official version, for Bibliography, the following notes have to be added for the standards indicated:

IEC 61850 (all parts)	NOTE	Harmonized as EN 61850 (all parts).
IEC 61850-8-1	NOTE	Harmonized as EN 61850-8-1 (not modified).
IEC 61850-9-2	NOTE	Harmonized as EN 61850-9-2 (not modified).
IEC 62439-6	NOTE	Harmonized as EN 62439-6 (not modified).
IEC 62439-7	NOTE	Harmonized as EN 62439-7 (not modified).

Annex ZA (normative)

Normative references to international publications with their corresponding European publications

The following documents are referred to in the text in such a way that some or all of their content constitutes requirements of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

NOTE 1 Where an International Publication has been modified by common modifications, indicated by (mod), the relevant EN/HD applies.

NOTE 2 Up-to-date information on the latest versions of the European Standards listed in this annex is available here: www.cenelec.eu.

<u>Publication</u>	<u>Year</u>	<u>Title</u>	<u>EN/HD</u>	<u>Year</u>
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 62439-1	-	Industrial communication networks - High availability automation networks -- Part 1: General concepts and calculation methods	EN 62439-1	-
IEC TR 61850-90-4	2013	Communication networks and systems for power utility automation - Part 90-4: Network engineering guidelines	-	-
ISO/IEC/IEEE 8802-3	2014	Standard for Ethernet	-	-
IEC/IEEE 61850-9-3	-	Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation	-	-
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 - Bridges and Bridged Networks	-	-
IETF RFC 2578	-	Structure of Management Information Version 2 (SMIv2), April 1999, http://tools.ietf.org/html/rfc2578	-	-
IETF RFC 3418	-	Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)	-	-

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

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.

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 $t_{\text{aliasMin}} = 13 \text{ s}$.




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 `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 `NodeRebootInterval` (see Table 8). `NodeRebootInterval` holds for all nodes in the network and therefore it is independent from the link speed.

Therefore, all nodes flush their entries before `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.

NOTE 2 Reception of a well-formed RCT is not a sufficient criterion to declare its source as DANP, since some protocols reply with the same frame as received.

4.3 PRP_Supervision frame

4.3.1 PRP_Supervision frame format

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																
6						msb	U/L	0								
8																
10																
12																
14																
16																
18																
20																
22																
24						msb	U/L	0								
26																
28																
30																
32						msb	U/L	0								
34																
36																
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.



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, 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.




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 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.

rows in Figure 16) are
 created from outside the
 Documents
 Document Sharing Hub

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: