

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 © 2010 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
Email: inmail@iec.ch
Web: 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.

- Catalogue of IEC publications: www.iec.ch/searchpub

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 It also gives information on projects, withdrawn and replaced publications.

- IEC Just Published: www.iec.ch/online_news/justpub

Stay up to date on all new IEC publications. Just Published details twice a month all new publications released. Available on-line and also by email.

- Electropedia: www.electropedia.org

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

- Customer Service Centre: www.iec.ch/webstore/custserv

If you wish to give us your feedback on this publication or need further assistance, please visit the Customer Service Centre FAQ or contact us:

Email: csc@iec.ch
Tel.: +41 22 919 02 11
Fax: +41 22 919 03 00

INTERNATIONAL STANDARD



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

INTERNATIONAL
ELECTROTECHNICAL
COMMISSION

PRICE CODE **XA**

ICS 25.040, 35.040

ISBN 978-2-88910-706-3

CONTENTS

FOREWORD.....	5
INTRODUCTION.....	7
1 Scope.....	8
2 Normative references.....	8
3 Terms, definitions, abbreviations, acronyms, and conventions.....	8
3.1 Terms and definitions.....	8
3.2 Abbreviations and acronyms.....	9
3.3 Conventions.....	9
4 Parallel Redundancy Protocol (PRP).....	9
4.1 PRP principle of operation.....	9
4.1.1 PRP network topology.....	9
4.1.2 PRP LANs with linear or bus topology.....	10
4.1.3 PRP LANs with ring topology.....	11
4.1.4 DANP node structure.....	11
4.1.5 PRP attachment of singly attached nodes.....	12
4.1.6 Compatibility between singly and doubly attached nodes.....	12
4.1.7 Network management.....	12
4.1.8 Implication on configuration.....	13
4.1.9 Transition to non-redundant networks.....	13
4.1.10 Duplicate handling.....	14
4.1.11 Configuration check.....	18
4.1.12 Network supervision.....	18
4.1.13 Redundancy management interface.....	19
4.2 PRP protocol specifications.....	19
4.2.1 Installation, configuration and repair guidelines.....	19
4.2.2 MAC addresses.....	20
4.2.3 Multicast MAC addresses.....	20
4.2.4 IP addresses.....	20
4.2.5 Nodes.....	20
4.2.6 Duplicate accept mode.....	21
4.2.7 Duplicate discard mode.....	21
4.3 PRP service specification.....	27
4.3.1 Arguments.....	27
4.3.2 NodesTable.....	28
4.3.3 PRP write.....	29
4.3.4 PRP read.....	30
5 High-availability Seamless Redundancy (HSR).....	31
5.1 HSR objectives.....	31
5.2 HSR principle of operation.....	31
5.2.1 Basic operation with a ring topology.....	31
5.2.2 DANH node structure.....	33
5.2.3 Topology.....	34
5.2.4 RedBox structure.....	40
5.3 HSR node specifications.....	42
5.3.1 Host sequence number.....	42
5.3.2 DANH receiving from its link layer interface.....	42

5.3.3	DANH receiving from an HSR port	42
5.3.4	DANH forwarding rules	43
5.3.5	CoS	44
5.3.6	Clock synchronization	44
5.3.7	Deterministic medium access	44
5.4	HSR RedBox specifications	44
5.4.1	RedBox properties	44
5.4.2	RedBox receiving from interlink	45
5.4.3	RedBox forwarding on the ring	46
5.4.4	RedBox receiving from an HSR port	46
5.4.5	Redbox proxy node table handling	47
5.4.6	RedBox CoS	47
5.4.7	RedBox clock synchronization	47
5.4.8	RedBox medium access	47
5.5	QuadBox specification	47
5.6	Association definition	47
5.7	Frame format for HSR	47
5.7.1	HSR-tagged frame format	47
5.7.2	HSR_Supervision frame	48
5.7.3	Constants	50
6	Protocol Implementation Conformance Statement (PICS)	51
7	PRP/HSR Management Information Base (MIB)	51
	Bibliography	58
	Figure 1 – PRP example of general redundant network	10
	Figure 2 – PRP example of redundant network as two LANs (bus topology)	10
	Figure 3 – PRP example of redundant ring with SANs and DANPs	11
	Figure 4 – PRP with two DANPs communicating	11
	Figure 5 – PRP RedBox, transition from single to double LAN	13
	Figure 6 – PRP frame extended by an RCT	15
	Figure 7 – PRP VLAN-tagged frame extended by an RCT	15
	Figure 8 – PRP constructed, padded frame closed by an RCT	16
	Figure 9 – PRP drop window on LAN_A	17
	Figure 10 – PRP drop window reduction after a discard	17
	Figure 11 – PRP frame from LAN_B was not discarded	18
	Figure 12 – PRP synchronized LANs	18
	Figure 13 – HSR example of ring configuration for multicast traffic	32
	Figure 14 – HSR example of ring configuration for unicast traffic	33
	Figure 15 – HSR structure of a DANH	34
	Figure 16 – HSR example of topology using two independent networks	35
	Figure 17 – HSR example of peer coupling of two rings	36
	Figure 18 – HSR example of connected rings	37
	Figure 19 – HSR example of coupling two redundant PRP LANs to a ring	38
	Figure 20 – HSR example of coupling from a ring node to redundant PRP LANs	39
	Figure 21 – HSR example of meshed topology	40
	Figure 22 – HSR structure of a RedBox	41

Figure 23 – HSR frame without VLAN tag 48

Figure 24 – HSR frame with VLAN tag 48

Table 1 – PRP_Supervision frame with VLAN tag 25

Table 2 – PRP constants 27

Table 3 – PRP arguments 28

Table 4 – PRP arguments 29

Table 5 – PRP write 29

Table 6 – PRP read 30

Table 7 – HSR_Supervision frame with optional VLAN tag 49

Table 8 – HSR Constants 50

IECNORM.COM : Click to view the full PDF of IEC 62439-3:2010

Withdrawing

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

This standard cancels and replaces IEC 62439 published in 2008. This first edition constitutes a technical revision.

This edition includes the following significant technical changes with respect to IEC 62439 (2008):

- adding a calculation method for RSTP (rapid spanning tree protocol, IEEE 802.1Q),
- adding two new redundancy protocols: HSR (High-availability Seamless Redundancy) and DRP (Distributed Redundancy Protocol),
- moving former Clauses 1 to 4 (introduction, definitions, general aspects) and the Annexes (taxonomy, availability calculation) to IEC 62439-1, which serves now as a base for the other documents,
- moving Clause 5 (MRP) to IEC 62439-2 with minor editorial changes,

- moving Clause 6 (PRP) was to IEC 62439-3 with minor editorial changes,
- moving Clause 7 (CRP) was to IEC 62439-4 with minor editorial changes, and
- moving Clause 8 (BRP) was to IEC 62439-5 with minor editorial changes,
- adding a method to calculate the maximum recovery time of RSTP in a restricted configuration (ring) to IEC 62439-1 as Clause 8,
- adding specifications of the HSR (High-availability Seamless Redundancy) protocol, which shares the principles of PRP to IEC 62439-3 as Clause 5, and
- introducing the DRP protocol as IEC 62439-6.

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

FDIS	Report on voting
65C/583/FDIS	65C/589/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 International Standard is to be read in conjunction with IEC 62439-1:2010, *Industrial communication networks – High availability automation networks – Part 1: General concepts and calculation methods*.

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

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

The committee has decided that the contents of this amendment and the base 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.

A bilingual version of this standard may be issued at a later date.

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

The IEC 62439 series specifies relevant principles for high availability networks that meet the requirements for industrial automation networks.

In the fault-free state of the network, the protocols of the IEC 62439 series provide ISO/IEC 8802-3 (IEEE 802.3) compatible, reliable data communication, and preserve determinism of real-time data communication. In cases of fault, removal, and insertion of a component, they provide deterministic recovery times.

These protocols retain fully the typical Ethernet communication capabilities as used in the office world, so that the software involved remains applicable.

The market is in need of several network solutions, each with different performance characteristics and functional capabilities, matching diverse application requirements. These solutions support different redundancy topologies and mechanisms which are introduced in IEC 62439-1 and specified in the other Parts of the IEC 62439 series. IEC 62439-1 also distinguishes between the different solutions, giving guidance to the user.

The IEC 62439 series follows the general structure and terms of IEC 61158 series.

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 detection of redundant frames given in 4.1.10.3, and concerning coupling of PRP and HSR LANs given in 5.4 (patent pending).

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 either free of charge or 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:

ABB Switzerland Ltd
Corporate Research
Segelhofstr 1K
5405 Baden
Switzerland

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://www.iec.ch/tctools/patent_decl.htm) 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 ISO/IEC 8802-3 (IEEE 802.3) (Ethernet) technology.

This part of the IEC 62439 series specifies two redundancy protocols based on the duplication of the LAN, resp. duplication of the transmitted information, designed to provide seamless recovery in case of single failure of an inter-switch link or switch in the network.

2 Normative references

The following referenced documents are indispensable for the application 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.

IEC 60050-191:1990, *International Electrotechnical Vocabulary – Chapter 191: Dependability and quality of service*

IEC 62439-1:2010, *Industrial communication networks – High availability automation networks – Part 1: General concepts and calculation methods*

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*

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

IEEE 802.1Q, *IEEE standards for local and metropolitan area network. Virtual bridged local area networks*

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, as well as in IEC 62439-1, apply, in addition to the following.

3.1.1

extended frame

frame that has been extended by a Redundancy Control Trailer

3.1.2

interlink

link that connects two network hierarchies

3.1.3

RedBox

device allowing to attach single attached nodes to a redundant network

3.1.4

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

3.1.5

HSR frame

frame that carries the HSR EtherType

3.2 Abbreviations and acronyms

For the purposes of this document, the following abbreviations and acronyms apply, in addition to those given in IEC 62439-1:

DANH	Double attached node implementing HSR
DANP	Double attached node implementing PRP
ICMP	Internet Control Message Protocol (part of the Internet protocol suite)
RCT	Redundancy Check Tag
SRP	Serial Redundancy Protocol
VDAN	Virtual Doubly Attached Node (SAN as visible through a RedBox)

3.3 Conventions

This document follows the conventions defined in IEC 62439-1.

4 Parallel Redundancy Protocol (PRP)

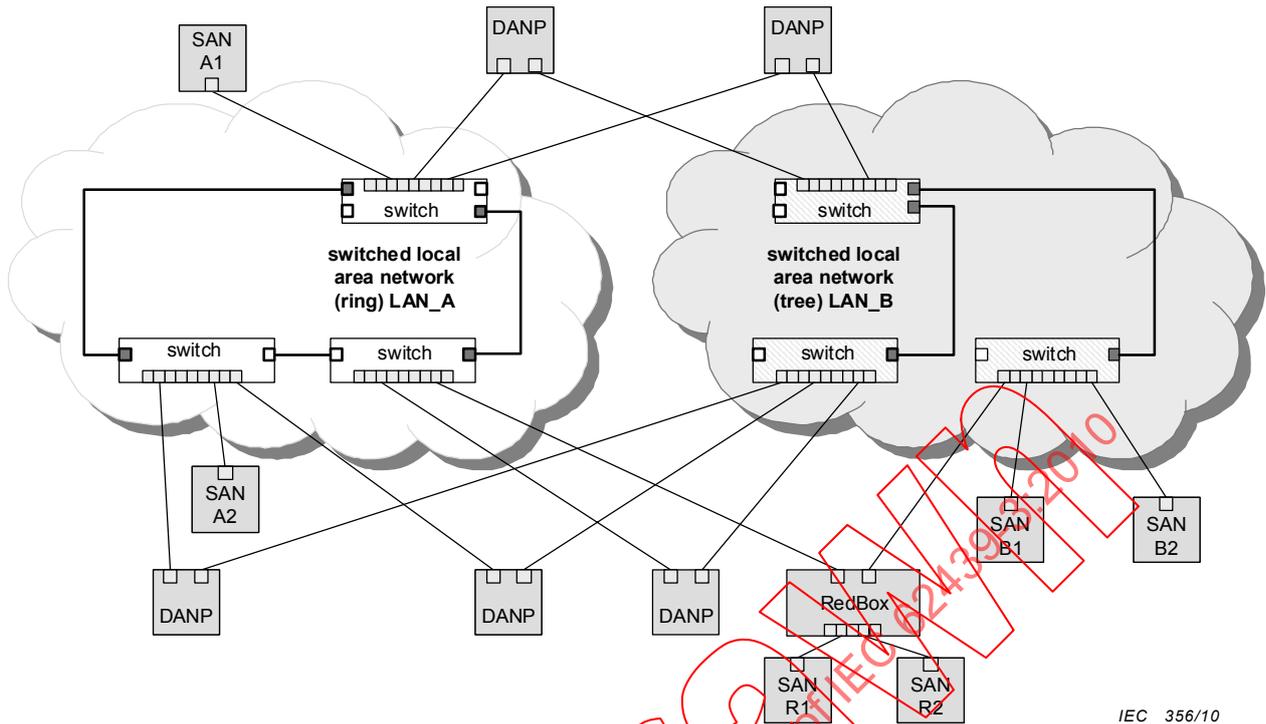
4.1 PRP principle of operation

4.1.1 PRP network topology

This redundancy protocol implements redundancy in the devices, through doubly attached nodes operating according to PRP (DANPs).

A DANP is attached to two independent 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 switched LANs, which can have any topology, e.g. tree, ring or meshed.



IEC 356/10

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.

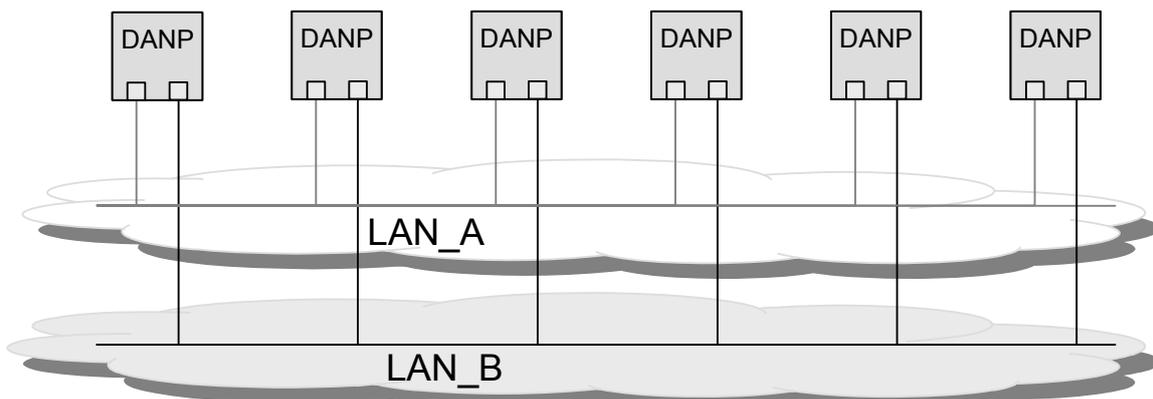
The two LANs follow configuration rules that allow the network management protocols such as Address Resolution Protocol (ARP) to operate correctly.

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 configuration,

Figure 2 draws a PRP network as two LANs in linear topology, which may also be a bus topology.



IEC 357/10

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.

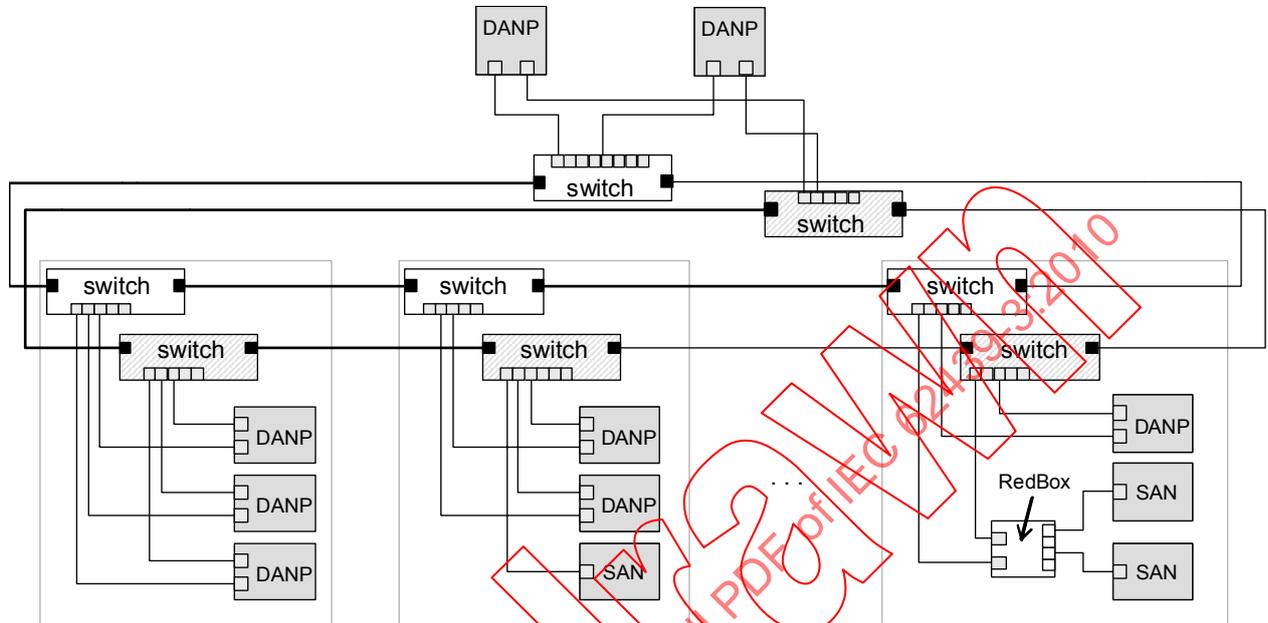


Figure 3 – PRP example of redundant ring with SANs and DANPs

IEC 358/10

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.

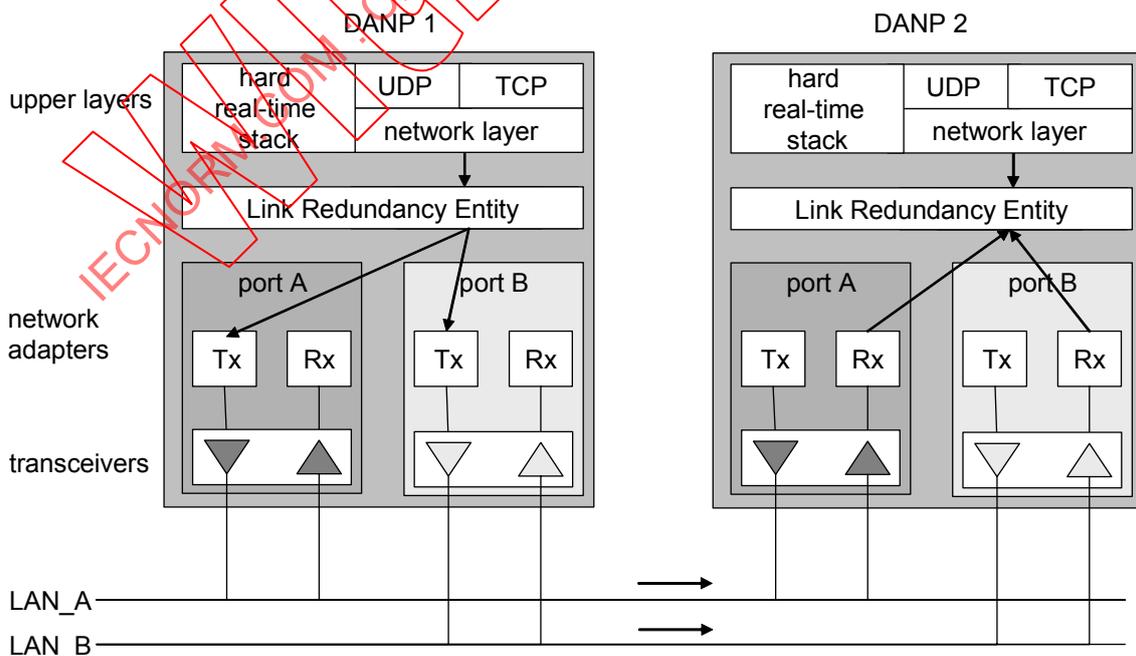


Figure 4 – PRP with two DANPs communicating

IEC 359/10

The Link Redundancy Entity (LRE) has two tasks: handling of duplicates and management of redundancy. This layer presents toward its upper layers the same interface as the network adapter of a non-redundant adapter.

When receiving a frame from the node's upper layers, the LRE sends the frame through both its ports at nearly the same time.

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

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

For management of redundancy, the LRE can append a Redundancy Check Trailer (RCT) including a sequence number to the frames it sends to keep track of duplicates. In addition, the LRE periodically sends PRP_Supervision frames and evaluates the PRP_Supervision frames of the 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. 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 with all DANPs.
- SANs can be attached over a RedBox (redundancy box) to both LANs, as Figure 1 shows for R1 and R2 (see also 4.1.9). Such SANs can communicate with all SANs, for instance SAN A1 and SAN R1 can communicate.

NOTE SANs do not need to be aware of PRP, they can be off-the-shelf computers.

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 switches 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 switch element according to IEEE 802.1D. For instance, the switch element may implement the MRP protocol, 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 International Standard. The supported mode is specified in the PICS (see 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. A SAN connected to one LAN cannot communicate directly to a SAN connected to the other LAN. Switches are always SANs. These SANs are not aware of PRP redundancy, so DANPs generate a traffic that these SANs understand. The condition is however that the SANs ignore the RCT in the frames, which should be the case since a SAN cannot distinguish the RCT from ISO/IEC 8802-3 (IEEE 802.3) padding. Conversely, DANPs understand the traffic generated by SANs, since these do not append a RCT. They only forward one frame to their upper layers since the SAN traffic uses one LAN only. If a DANP cannot positively identify that the remote device is a DANP, it considers it as a SAN.

4.1.7 Network management

A node has the same MAC address on both ports, and only one set of IP addresses assigned to that address. This makes redundancy transparent to the upper layers. Especially, this allows the Address Resolution Protocol (ARP) to work the same as with a SAN. Switches in a LAN are not doubly attached devices, and therefore all managed switches have different IP addresses. A network management tool is preferably a DANP and can access nodes and

switches as if they all belong to the same network. Especially, network management implemented in a DANP is able to see SANs connected to either LAN.

Some applications require different MAC addresses on the redundant ports, and these MAC addresses may be different from the default MAC address of that node. This involves address substitution mechanisms which are not specified in this International Standard. However, the basic protocol and the frame format are prepared for such extension. Nodes that support MAC address substitution are indicated as supporting PICS_SUBS.

4.1.8 Implication on configuration

Since the same frame can come from the two ports with significant time difference, the period of cyclic time-critical data must be chosen so that it considers the difference between worst case and best case path latency between publisher and subscriber.

4.1.9 Transition to non-redundant networks

The mechanism of duplicate rejection can be implemented by the 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 VDA or virtual DANs) and multicasts supervision frames on their behalf, appending its own information. The RedBox is itself a DANP and has its own IP address for management purposes, but it may also perform application functions.

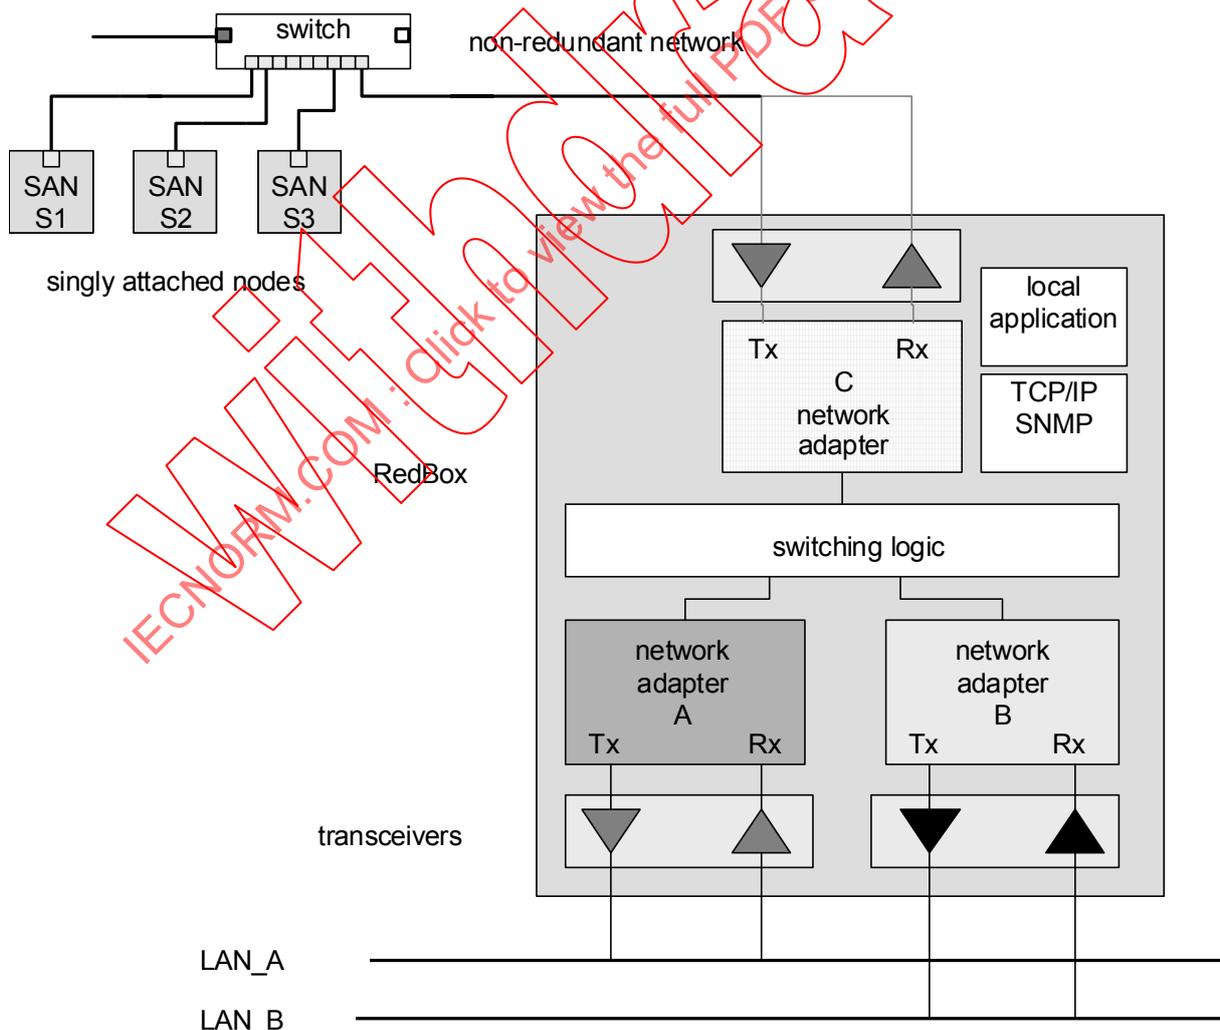


Figure 5 – PRP RedBox, transition from single to double LAN

4.1.10 Duplicate handling

4.1.10.1 Methods for handling duplicates

Since a DANP receives the same frame over both adapters, when both are operational, it should keep one and ignore the duplicate.

There are two methods for handling duplicates:

- a) duplicate accept, in which the sender LRE uses the original frames and the receiver LRE forwards both frames it receives to its upper protocol layers;
- b) duplicate discard, in which the sender LRE appends a redundancy control trailer to both frames it sends and the receiver LRE uses that redundancy control trailer to send only the first frame of a pair to its upper layers and filter out duplicates.

4.1.10.2 Duplicate accept

This method does not attempt to discard duplicates at the link layer. The sender LRE sends the same frame as it would in the non-redundant case over both LANs. The receiver's LRE forwards both frames of a pair (if both arrive) to its upper layers, assuming that well-designed network protocols and applications are able to withstand duplicates – indeed IEEE 802.1D explicitly states that it cannot ensure freedom of duplicates.

The internet stack, consisting of a network layer with an UDP and a TCP transport layer, is assumed to be resilient against duplicates. The TCP protocol is designed to reject duplicates, so it discards the second frame of a pair. The UDP layer is by definition connectionless and unacknowledged. All applications that use UDP are assumed to be capable of handling duplicates, since duplication of frames can occur in any network. In particular, a UDP frame is assumed to be idempotent, i.e. sending it twice has the same effect as sending it once. Administrative protocols of the internet such as ICMP and ARP are not affected by duplicates, since they have their own sequence numbering.

Real-time stack that operate on the publisher-subscriber principle are not affected by duplicates, since only the latest value is kept. Duplicate reception increases robustness since a sample that gets lost on one LAN is usually received from the other LAN.

Therefore, one can assume that handling of duplicates is taken care of by the usual network protocols, but one has to check if each application complies with these assumptions.

This simple duplicate accept method does not provide easy redundancy supervision, since it does not keep track of correct reception of both frames. The receiver would need hash tables to know that a frame is the first of a pair of a duplicate, and could for this effect store the CRC and length of each frame as a hash code. Such redundancy supervision method is however not specified in this International Standard, but it is not excluded.

4.1.10.3 Duplicate discard in the link layer

4.1.10.3.1 Principle

It is advantageous to discard duplicates already at the link layer.

Without duplicate discard, the processor receives twice as many interrupt requests as when only one LAN is connected. To offload the application processor, the LRE can perform Duplicate Discard, possibly with an independent pre-processor or an intelligent Ethernet controller. This allows at the same time to improve the redundancy supervision.

The duplicate discard protocol uses an additional four-octet field in the frame, the Redundancy Control Trailer (RCT), which the LRE inserts into each frame that it receives from the upper layers before sending, as Figure 6 shows. The RCT consists of the following parameters:

- a) 16-bit sequence number (SequenceNr);
- b) 4-bit LAN identifier (Lan);
- c) 12 bit frame size (LSDU_size).

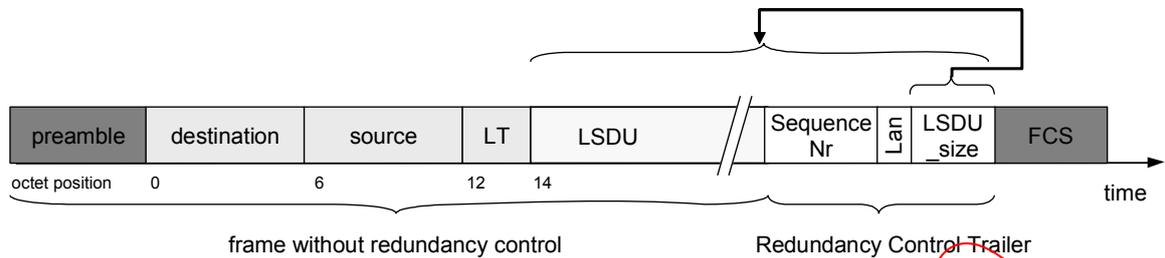


Figure 6 – PRP frame extended by an RCT

IEC 361/10

4.1.10.3.2 Use of SequenceNr

Each time a LRE sends a frame to a particular destination, it increases the sequence number corresponding to that destination and sends both (nearly identical) frames over both LANs.

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

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 same sequence number and the same size. However, SANs send on one LAN only, and the source will not be the same as that of another frame, so a frame from a SAN will never be discarded.

4.1.10.3.3 Use of LAN

The field LAN can take one of two values: 1 010 indicating that the frame has been sent over LAN_A and 1 011 indicating that the frame has been sent over LAN_B. This allows detecting installation errors.

4.1.10.3.4 Use of LSDU_size

To allow the receiver LRE to distinguish easily frames coming from nodes that obey to the PRP from the non-redundant ones, the sender LRE appends to the frame the length of the link service data unit (LSDU) in octets in a 12-bit field.

EXAMPLE If the frame carries a 100-octets LSDU, the size field equals LSDU+RCT: $104 = 100 + 4$.

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

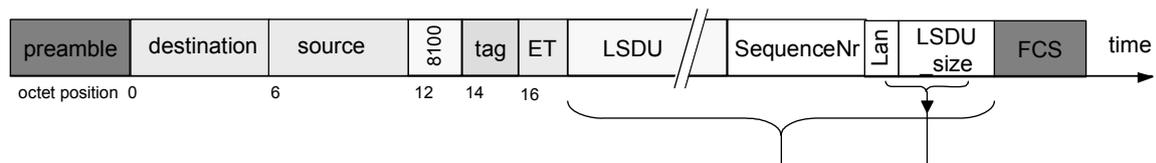


Figure 7 – PRP VLAN-tagged frame extended by an RCT

IEC 362/10

The receiver scans the frames, preferably starting from the end. If it detects that the 12 bits before the end correspond to the LSDU size, and that the LAN identifier matches the identifier of the LAN it is attached to (see 4.1.11), the frame is a candidate for rejection.

Since short frames need padding to meet the minimum frame size of 64 octets, the sender already includes the padding to speed up scanning from behind, as Figure 8 shows.

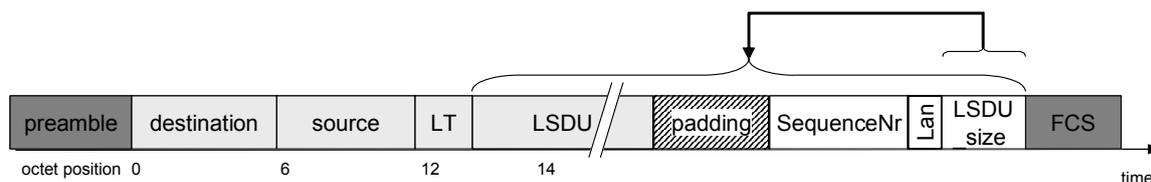


Figure 8 – PRP constructed, padded frame closed by an RCT

IEC 363/10

NOTE A VLAN-tagged frame can pass several switches which may remove or insert VLAN tags. If the sender observes the ISO/IEC 8802-3 (IEEE 802.3) rule to send a minimum frame size of 68 octets for a VLAN-tagged frame and of 64 for a VLAN-untagged frame, there should never be a situation in which there is padding before and after the RCT. Scanning from behind is specified as a matter of precaution.

4.1.10.3.5 Frame size restriction

Appending the RCT could generate oversize frames that exceed the `maxValidSize` foreseen by ISO/IEC 8802-3 (IEEE 802.3).

To maintain compliance with IEEE 802.3:2005, the communication software in a DANP using duplicate discard is configured for a maximum payload size of 1 496 octets.

NOTE Longer payloads would work in most cases, but this requires previous testing. Many switches are dimensioned for double-VLAN-tagged (non-IEEE 802.3 compliant) frames that have a maximum size of 1 526 octets. Most Ethernet controllers are certified up to 1 528 octets. Most switches would forward correctly frames of up to 1 536 octets, but this cannot be relied upon.

4.1.10.3.6 Discard algorithm

The following algorithm is optional, other methods such as hash table can be used.

The receiver assumes that frames coming from a DANP are sent in sequence with increasing sequence numbers. The sequence number expected for the next frame is kept in the variables `ExpectedSeqA`, respectively `ExpectedSeqB`.

At reception, the correct sequence can be checked by comparing `ExpectedSeqA` with the received sequence number in the RCT, `CurrentSeqA`. Regardless of the result, `ExpectedSeqA` is set to one more than `CurrentSeqA` to allow checking the next expected sequence number on that line. The same applies to `ExpectedSeqB` and `CurrentSeqB` on LAN_B.

Both LANs thus maintain a sliding drop window of contiguous sequence numbers, the upper bound being `ExpectedSeqA` (the next expected sequence number on that LAN), excluding that value, the lower bound being `StartSeqA` (the lowest sequence number that leads to a discard on that LAN) as Figure 9 shows for LAN_A. The same applies to `ExpectedSeqB` and `StartSeqB` on LAN_B.

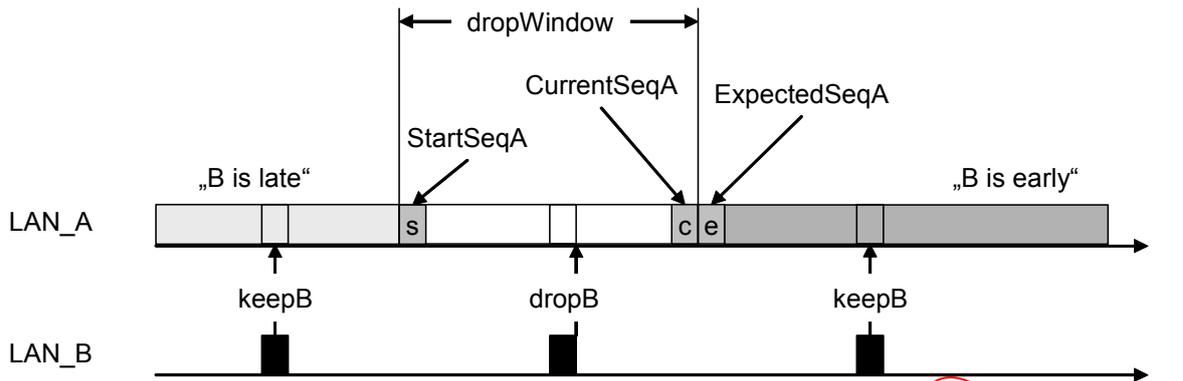


Figure 9 – PRP drop window on LAN_A

IEC 364/10

After checking the correct sequence number, the receiver decides whether to discard the frame or not. Assuming that LAN_A has established a non-void drop window (as in Figure 9), a frame from LAN_B whose sequence number CurrentSeqB fits into the drop window of A is discarded (dropB in Figure 9). In all other cases, the frame is kept and forwarded to the upper protocol layers (keepB in Figure 9).

Discarding the frame (dropB in Figure 9) shrinks the drop window size on LAN_A since no more frames from B with an earlier sequence number are expected, thus StartSeqA is increased to one more than the received CurrentSeqB. Also, the drop window on B is reset to a size of 0 (StartSeqB = ExpectedSeqB), since obviously B lags behind A and no frames from A should be discarded, as Figure 10 shows.

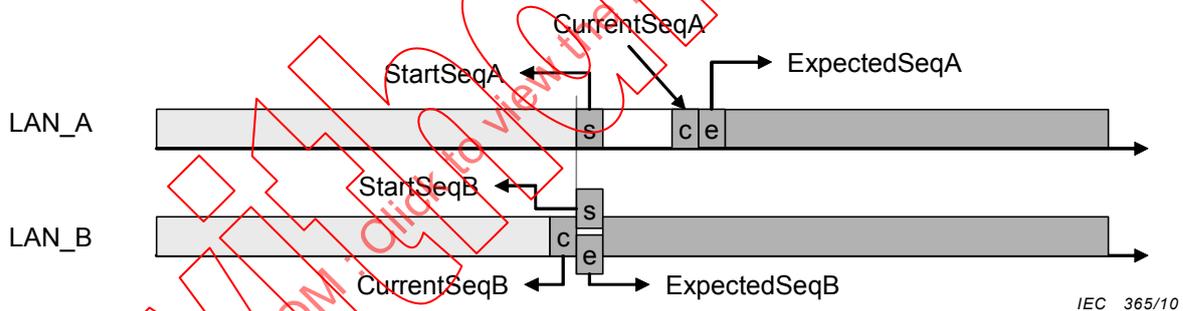


Figure 10 – PRP drop window reduction after a discard

IEC 365/10

In the situation of Figure 10, if several frames come in sequence over the same LAN_A, but none on LAN_B, they are kept since their CurrentSeqA is outside the drop window of LAN_B, and the drop window of LAN_A grows by one position. If frames keep on coming over LAN_A but not LAN_B when the maximum drop window size is reached, StartSeqA is also incremented to slide the drop window.

When a received frame is out of the drop window of the other LAN, it is kept and the drop window of that line is reduced to a size of 1, meaning that only a frame from the other line with the same sequence number is discarded, while the drop window of the other line is reset to 0, meaning that no frame is discarded, as Figure 11 shows.

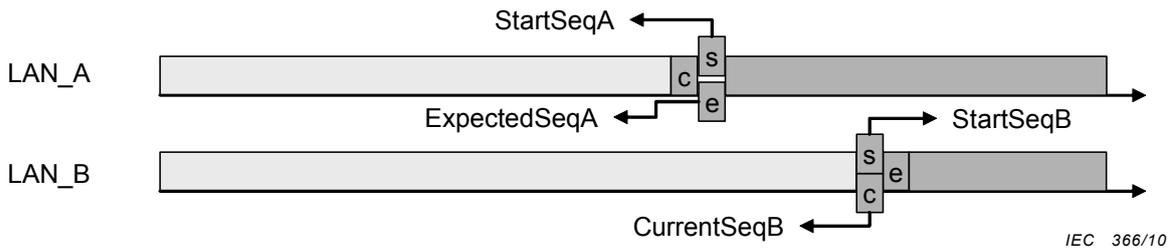


Figure 11 – PRP frame from LAN_B was not discarded

The most common situation is when the two lines are synchronized and both drop windows are reduced to 0, meaning that the first frame to come next is kept and the drop window is opened by one to allow only a frame with the same sequence number as the one already received, as Figure 12 shows.

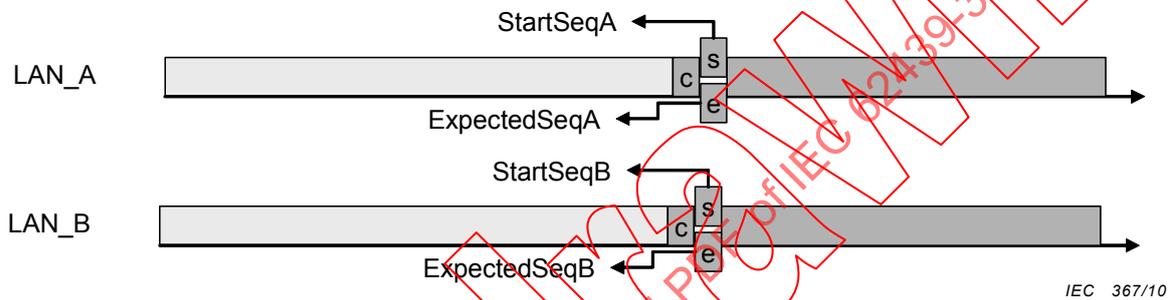


Figure 12 – PRP synchronized LANs

The sequence counter has 16 bits, which allows a drop window size of 32 768, a size large enough so that even under the worst case network delays and highest frame rate the sequence numbers do not wrap-around.

There is no change to this algorithm when frames come out of sequence.

This method can be defeated by some situations, for instance nodes failing and recovering or reconnection of a damaged LAN after a long time, but in case of doubt, duplicates are accepted so that no frame is lost.

Annex A discloses a pseudo-code for the duplicate discard algorithm.

4.1.11 Configuration check

The remaining 4 bits of the RCT carry a distinct identifier for LAN_A or LAN_B, specifically the codes 1 010 (“A”) and 1 011 (“B”). Therefore, the frames 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 length information in its last 12 bits, but it increments the error counters CntErrWrongLanA or CntErrWrongLanB since this could hint at a configuration error. Since this kind of error is permanent, it is detected rapidly.

4.1.12 Network supervision

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

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

To this effect, all senders and receivers maintain tables of nodes with which they communicate 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.

At the same time, these tables allow to establish connections to synchronize the sequence numbers and detect sequence gaps and missing nodes.

Since the protocol is loosely connection oriented, the sequence numbers corresponding to non-existent nodes are cleaned up by a low priority task after a time NodeForgetTime.

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.

4.1.13 Redundancy management interface

Redundant devices and links are useless without network management supervising this 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 adapter (each LAN) a counter of received messages and of messages received with an error.

The LAN statuses appear as SNMPv1 or SNMPv2/v3 variables. This allows using the same tools for managing the nodes and the switches.

NOTE SNMP is part of the IP protocol suite.

4.2 PRP protocol specifications

4.2.1 Installation, configuration and repair guidelines

NOTE These guidelines are to be followed at installation time, they do not apply to conformance testing of the devices.

4.2.1.1 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.

4.2.1.2 Labelling cables

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

4.2.1.3 Labelling switches

Switches in the two LANs shall have a distinct label or colour for each A or B.

4.2.1.4 Independent operation

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

4.2.1.5 Configuration

All DANPs shall be configured with the same multicast address for PRP_Supervision frames.

All DANPs shall be configured with the same LifeCheckInterval.

4.2.2 MAC addresses

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

This 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).

If a DANP implements PICS_SUBS, the MAC address shall be the MAC address of adapter A and adapter B may use a different MAC address, which shall be unique within the whole network (LAN_A plus LAN_B).

NOTE Nodes supporting PICS_SUBS are expected to behave as a DANP that has the default MAC address. Address substitution is not specified in this International Standard.

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 4.2.7.6.

4.2.4 IP addresses

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

NOTE A device may have several IP addresses.

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

Switches on LAN_A and LAN_B are considered as SANs and shall have different IP addresses for the purpose of network management.

4.2.5 Nodes

4.2.5.1 Node types

Doubly attached nodes according to the parallel redundancy protocol (DANP) shall have two network adapters (adapter A and adapter B) that have the same abilities, and in particular could be used alternatively if only one LAN is connected, adapter A being connected to LAN_A and adapter B to LAN_B.

Singly Attached Nodes (SAN) have only one adapter for the purpose of this protocol and may be attached to either LAN.

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

4.2.5.2 Labelling connectors

This subclause 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 shall be the upper connector and LAN_B the lower connector in its normal position.

When connectors are ordered horizontally, the left connector shall 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

4.2.6.1 Sending

The sender shall send the frame it receives from its upper layers unchanged over both its adapters so that the two frames appear on the respective LANs.

4.2.6.2 Receiving

The receiver shall forward frames received from both adapters to the upper layers.

NOTE This specification is only testable indirectly, by counting the number of frames over the MIB.

4.2.7 Duplicate discard mode

4.2.7.1 Nodes table

A node shall maintain a table with an entry for each node (SAN or DANP) to which it sends a frame, or from which it receives a frame, using the MAC address as a key. The table shall contain the following information for each unicast, multicast or broadcast address sent by that node:

- a) **SendSeq**
a 16-bit sequence number used by this node for sending to that remote node or multicast or broadcast address (wrapping through zero)
- b) **ExpectedSeqA** and **ExpectedSeqB**
for each adapter A and B, a 16-bit sequence number indicating the sequence number used last by the remote node to communicate with this node on that LAN, incremented by one (wrapping through zero)
- c) **CntErrOutOfSequenceA** and **CntErrOutOfSequenceB**
for each adapter A and B, a 32-bit error counter indicating that a frame from the remote node was not received in sequence over that LAN
- d) **StartSeqA** and **StartSeqB**
for each adapter A and B, a 16-bit cursor that limits the drop window
- e) **CntReceivedA** and **CntReceivedB**
for each adapter A and B, a 32-bit counter indicating the number of frames received over the adapter
- f) **CntErrWrongLanA** and **CntErrWrongLanB**
for each adapter A and B, a 32-bit counter indicating the number of mismatches on each adapter
- g) **TimeLastSeenA** and **TimeLastSeenB**
for each adapter A and B, a time field indicating when this node received last a frame from the remote node. This field is in some cases updated at sending to keep track of ageing.
- h) **SanA** and **SanB**
for each adapter A and B, a boolean indicating that the remote node is probably a SAN and/or that the remote node uses duplicate accept (see 4.2.7.4.2).

NOTE 1 The table contains for each remote node one row for the unicast frames and one row for each multicast or broadcast address that remote node is sending. It contains one row for each unicast, multicast or broadcast address this node is sending.

NOTE 2 Some fields are irrelevant for a SAN.

NOTE 3 This is a conceptual view, distinct tables for destination and source nodes could be implemented.

4.2.7.2 Redundancy Control Trailer (RCT)

The Redundancy Control Trailer (RCT) inserted into each DANP frame shall consist of four octets, structured in the following way (in the order of transmission):

- a) a 16-bit sequence number (SequenceNr) transmitted with the most significant 8 bits in the first octet, which reflects the counter SendSeq of the nodes table for the destination of the frame (see 4.2.7.1).
- b) a 4-bit LAN identifier (Lan) transmitted as the most significant 4 bits of the third octet, which carries the sequence “1010” for LAN_A, respectively the sequence “1011” for LAN_B.
- c) a 12 bit LSDU size (LSDU_size) whose most significant 4 bits are transmitted in the least significant 4 bits of the 3rd 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 8802-3 (IEEE 802.3) and IEEE 802.1Q (octet offset 12-13 without LAN header or 16-17 with VLAN header) to the RCT, excluding the PT, and the frame part after the RCT, but including the RCT itself.

NOTE Padding inserted before the RCT is included in the LSDU size, padding inserted after the RCT is not included in the LSDU size.

4.2.7.3 Sending (duplicate discard mode)

4.2.7.3.1 Frame size control

The sender shall have the ability to limit the LSDU size so that the complete frame, including the four-octet redundancy control trailer, does not exceed the maximum size allowed on the LAN when it operates in the duplicate discard mode.

NOTE 1 This maximum size is currently 1 518 octets for VLAN untagged frames according to IEEE 802.3:2005.

NOTE 2 This specification does not apply to the LRE, but to its upper layers.

4.2.7.3.2 Sending and nodes table

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

- a) update the nodes table:
 - if the destination address (single cast, multicast or broadcast address) is not yet in the nodes table, create an entry in that table and record as TimeLastSeenA and TimeLastSeenB the current time. If the destination is a unicast address, set the SanA and the SanB to 1, if it is a multicast or broadcast address set them to 0. All other values shall be reset to 0, except for the sequence number SendSeq that may take an arbitrary value, preferably the value 1;
 - if the destination address (single cast, multicast or broadcast address) is already in the nodes table, increment the sequence number SendSeq for that address, wrapping over through 0;
 - if the destination address is a multicast address or the broadcast address, update in addition the TimeLastSeenA and TimeLastSeenB counters.

NOTE 1 Updating TimeLastSeenA, respectively TimeLastSeenB at sending initializes the ageing time for the remote node. The receiving process actualizes this time value when it receives a frame from that node. A time-out process removes the entry.

NOTE 2 Duplicate discard is assumed for multicast/broadcast addresses, since no PRP_Supervision frame tells the mode. For unicast addresses, the remote node is likely a SAN on LAN_A or LAN_B. If the destination is a DANP, an entry in the nodes table probably exists due to a previously received PRP_Supervision frame, or one is coming soon.

- b) send
 - if either SanA or SanB is set, send the frame unchanged over the corresponding adapter;
 - if both are set, send the frame unchanged over both adapters;

- if none is set, append the Redundancy Control Trailer (RCT) between the LSDU (payload) and before the FCS, preferably just before the FCS if padding is used and send the appended frame with LAN identifier “A” through its adapter A and the frame with LAN identifier “B” through its adapter B, under the same conditions as 4.2.6.1.

4.2.7.4 Receiver (duplicate discard mode)

4.2.7.4.1 Receiving and nodes table

On reception of a frame that is not a BPDU according to IEEE 802.1Q over either adapter, a node shall:

- a) if the adapter signals that the frame is in error, increment the error counter of the respective adapter CntErrorsA or CntErrorsB and ignore the frame;
- b) otherwise
 - if this frame is not a PRP_Supervision frame and not a BPDU and its source is not yet in the nodes table, create an entry in the nodes table for that source MAC address assuming it is a SANA or a SANB, depending which LAN the frame arrives on;
 - if the frame is received 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 this frame is a PRP_Supervision frame, and its source is not yet in the nodes table, create an entry in the nodes table for that source assuming DANP duplicate accept or duplicate discard according to the PRP_Supervision frame contents. If the source is already in the nodes table, update its 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 nodes table for that source;
 - increment by one (wrapping through 0) the counters CntReceivedA, respectively CntReceivedB of the nodes table for that source and address kind.

NOTE Updating SanA and SanB allows to move a 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.

4.2.7.4.2 Identification of frames associated with the duplicate discard mode

A receiver shall identify as a duplicate candidate a frame whose last 12 bits before the FCS match the physical size of the LSDU as defined in Figure 6, except for small frames that use padding, for which the receiver shall scan the frame backwards until it finds a matching size field, stopping when reaching the LT field.

NOTE 1 Small frames using padding are smaller than 64 octets.

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

4.2.7.4.3 LAN identification

A receiver shall check for a frame identified as a duplicate candidate that the four bits previous to the size are either 1 010 (A) or 1 011 (B).

A receiver shall increment the CntErrWrongLanA, respective CntErrWrongLanB counter of the source device in the nodes table if the LAN identifier does not match the adapter from which it received the frame and forward the unchanged frame to its upper layers.

NOTE If one SAN is moved from LAN_A to LAN_B, it will first be considered DANP duplicate accept for the duration of NodeForgetTime before it becomes a SAN B.

4.2.7.4.4 Duplicate discarding

A receiver can use any method to discard duplicates, provided that this method does not discard a frame sent as single or both frames of a pair, while it is permitted that in case of doubt, both frames of a pair can be passed to the higher protocol layers.

The following drop window algorithm is recommended and uses the following fields: source MAC address, destination MAC address (or multicast address), RCT.

4.2.7.4.5 Drop window

A receiver shall consider for each LAN and source node the drop window as the range of sequence numbers from StartSeqA to (excluded) ExpectedSeqA, respectively StartSeqB to (excluded) ExpectedSeqB, in Modulo 16 arithmetic.

4.2.7.4.6 Sequence check

The receiver shall check if the received frame is in sequence by comparing it with the ExpectedSeqA, respectively ExpectedSeqB of the LAN over which it was received, and increment the error counter CntErrOutOfSequenceA, respectively CntErrOutOfSequenceB if they are not equal, and then increment ExpectedSeqA, respectively ExpectedSeqB.

4.2.7.4.7 Frame discard

If the sequence number of a frame that is a duplicate candidate is within the drop window of the other LAN, the receiver shall discard that frame, reset the drop window of the LAN over which the frame was received to 0 (StartSeqA := ExpectedSeqB respectively StartSeqB := ExpectedSeqA) and move the lower bound of the drop window on the other LAN to one position ahead of the received frame (StartSeqA := StartSeqB).

4.2.7.4.8 Frame keeping

If the sequence number of a frame that is a duplicate candidate is outside the drop window of the other LAN, the receiver shall forward the frame to the upper layers.

If the sequence number is in sequence on that LAN, the receiver shall, if the maximum window size DropWindowMax (see 4.2.7.8) has been reached, increase by one the lower drop window bound for the LAN over which the frame was received, StartSeqA or StartSeqB.

If the received sequence number is out-of-sequence, the receiver shall reset the drop window on that LAN to one (e.g. StartSeqB := CurrentSeqB).

4.2.7.4.9 Transparent reception

If the configuration setting TransparentReception of the node is set, the receiver shall not remove the RCT before transferring the frame to the upper layers.

If the configuration setting TransparentReception of the node is not set, the receiver shall remove the RCT on frames where it has identified the presence of the RCT.

4.2.7.5 Cleanup of the nodes table

A node shall clear a nodes table entry when the time elapsed since reception of a frame from that source over both TimeLastSeenA and TimeLastSeenB exceeds NodeForgetTime (see 4.2.7.8).

NOTE It is sufficient to check the whole nodes table every NodeForgetTime for stale entries.

4.2.7.6 PRP_Supervision frame

4.2.7.6.1 Sending

Each DANP shall multicast a PRP_Supervision frame over both its adapters with the format specified in Table 1 every LifeCheckInterval (see 4.2.7.8). This format shall also be used when the node is operating in duplicate accept mode.

Table 1 – PRP_Supervision frame with 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	PRP_SourceAddress (MAC address of the adapter)															
10									lsb							
12	ptid (0x8100 for VLAN or 0x88FB for PRP)															
14	prio		cti		vlan_identifier											
16	pt (= 0x88FB for PRP)															
18	LID				PRP_Ver < 64											
20	PRP_TLV.Type = 20 or 21								PRP_TLV.Length = 12							
22						msb	U/L	0								
24	MacAddressA (MAC address A of the DANP)															
26									lsb							
28						msb	U/L	0								
30	MacAddressB (MAC address B of the DANP)															
32									lsb							
34	PRP_TLV2.Type = 30 or 31								PRP_TLV2.Length = 6							
36						msb	U/L	0								
38	RedBoxMacAddress															
40									lsb							
	Padding to 64 octets (no VLAN) or to 68 octets (VLAN)															
60	SequenceNr															
62	Lan (0x1010 or 0x1011)						LSDU_size = 46									
64	FCS															
66																

4.2.7.6.2 PRP_Supervision frame contents

PRP_DestinationAddress

Reserved multicast address 01-15-4E-00-01-XX shall be used for this protocol. By default XX is "00", but if conflicts arise, XX can be configured to take any value between 0x00 and 0xFF.

PRP_SourceAddress

MAC address of the sending adapter.

PRP_Ver

Indicates the protocol version, set to “0” (zero) for this version of 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. The version shall not exceed the value of 64, since the same beacon is used for HSR.

PRP_TLV.Type

Indicates the operation mode and shall have a value of 20 to indicate that the node supports the duplicate discard or a value of 21 to indicate that it implements duplicate accept. Other values are reserved.

PRP_TLV.Length

Indicates the length of the following MAC addresses (12).

MacAddressA and MacAddressB

MAC addresses used by each port. These addresses shall be identical except if address substitution (PICS=PRP_SUBS) is supported by the sender.

PRP_TLV2

This field shall be set to 0 if the source node is not a RedBox (see 4.2.7.6.3)

SequenceNr

Sequence number used for PRP_Supervision frames.

Lan

LAN over which this PRP_Supervision frame is sent.

LSDU_size

Size of the LSDU, always 46 (independently if VLAN tagging is used or not).

The following fields are only sent by a RedBox when it relays frames on behalf of a SAN and at least the next two octets shall be 0 for other nodes.

NOTE 1 Octets with offset 14 to 17 are inserted only if VLAN according to 802.1D is used.

NOTE 2 The frame has a size of 68 octets if VLAN-tagging is used to avoid padding if a switch removes the VLAN tag.

4.2.7.6.3 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 TLV field a second TLV field with the following contents:

PRP_TLV2.Type

Indicates the operation mode and shall have a value of 30 to indicate that the node is a RedBox or a value of 31 to indicate that it is a VDAN. Other values are reserved. This field shall only be sent by a RedBox, otherwise it shall be zero.

PRP_TLV2.Length

Indicates the length of the following MAC address (6 for a RedBox, 0 otherwise).

RedBoxMacAddress

MAC address of the RedBox that acts as proxy for the other device. This field shall only be sent by a RedBox, otherwise it shall be zero.

4.2.7.6.4 Reception of a PRP_Supervision frame

When receiving a PRP_Supervision frame over any LAN, a node shall create an entry in the nodes table corresponding to the MacAddressA 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 MacAddressA and MacAddressB are different, this indicates that the sending node supports PICS_SUBS. If the receiving node supports PICS_SUBS, a receiving node shall, in all frames it receives from that node over adapter A respective adapter B, substitute the received MAC

address by the default MAC address of that node (which may be identical to MacAddressA) before forwarding the frame to the upper layers.

4.2.7.6.5 Non-Reception of a PRP_Supervision 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 right 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 a SAN from a DANP in duplicate accept mode with one line disconnected.

4.2.7.7 Switching end node

If this setting is enabled, the node shall act as a switching end node for its two ports, implementing either:

- SRP (serial redundancy protocol), a subset of IEEE 802.1D Clause 8 in which its ports may only have the root or alternate/backup role, subject to PICS_SRP or
- RSTP, (rapid spanning tree protocol), the IEEE 802.1D Clause 8 in which its ports can take the root, alternate/backup or designated role, subject to the PICS PRP_RSTP or
- MRP, see IEC 62439-2, subject to the PICS PRP_MRP

NOTE 1 The switching end node setting supports attachment of a DANP to two switches of the same LAN to implement a partial redundancy topology. Activating this setting implies duplicate accept. 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 one of these modes is enabled.

4.2.7.8 Constants

The constant parameters are shown in Table 2.

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

Table 2 – PRP constants

Constant	Description	Default value
LifeCheckInterval	How often a node sends a PRP_Supervision frame	2 000 ms
NodeForgetTime	Time after which a node entry is cleared	60 000 ms
DropWindowMax	Max size of drop window	32 768 frames

4.3 PRP service specification

4.3.1 Arguments

These arguments are used in both the command and the response (see Table 3 below). In a command (PRP write), they indicate the desired setting and in a status (PRP read), they indicate the actual setting.

Table 3 – PRP arguments

Argument	Definition	Data type
NodeName	Node name in the LRE	VisibleString32
ManufacturerName	Name of the LRE manufacturer	VisibleString255 (can be read only)
VersionName	Version of the LRE software	VisibleString32
MacAddressA	MAC address to be used by network interface A	Unsigned48
MacAddressB	MAC address to be used by network interface B	Unsigned48
AdapterActiveA	Adapter A is commanded to be active or responds that it is active if true	Boolean1
AdapterActiveB	Adapter B is commanded to be active or responds that it is active if true	Boolean1
DuplicateDiscard	Duplicate discard algorithm is (to be) used at reception and the RCT is (to be) appended at sending if true	Boolean1
TransparentReception	RCT is not (to be) removed when forwarding to the upper layers if true	Boolean1
SwitchingEndNode	if 0: LRE is not (to be) configured as a switching node if 1: LRE is (to be) configured as an SRP switching node if 2: LRE is (to be) configured as an RSTP switching node if 4: LRE is (to be) configured as an MRP switching node if 5: LRE is (to be) configured as an HSR switching node if 6: LRE is (to be) configured as a RedBox in H mode if 7: LRE is (to be) configured as a RedBox in A mode if 8: LRE is (to be) configured as a RedBox in B mode	Integer8
NodesTableClear	Nodes table is (to be) cleared if true	Boolean1
SupervisionAddress	Address to be used for PRP_Supervision frames	Unsigned 48
LifeCheckInterval	Interval at which the PRP_Supervision frame is (to be) sent in milliseconds	Unsigned16
NodeForgetTime	Interval at which the nodes table entry of a node is (to be) cleared, in seconds	Unsigned16
DropWindowMax	Maximum size of the drop window to be used	Unsigned16
CntTotalSentA	Number of frames sent over adapter A	Unsigned32
CntTotalSentB	Number of frames sent over adapter B	Unsigned32
CntTotalReceivedA	Number of frames received over adapter A	Unsigned32
CntTotalReceivedB	Number of frames received over adapter B	Unsigned32
CntErrorsA	Number of transmission errors on adapter A, as signalled by the adapter	Unsigned32
CntErrorsB	Number of transmission errors on adapter B, as signalled by the adapter	Unsigned32
CntNodes	Number of nodes in NodesTable	Unsigned16
NodesTable	Records for all nodes that have been detected within the last NodeForgetTime the following fields	Sequence, see 4.3.2

4.3.2 NodesTable

The node table (Table 4) keeps for network management purpose the record of all nodes that have been detected on the network.

NOTE 1 The key attribute of the Nodes table is MacAddressA as received in the PRP_Supervision frame sent by a DANP.

NOTE 2 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 (See 4.2.7.1).

Table 4 – PRP arguments

Argument	Definition	Data Type
MacAddressA	MAC address of the source node (6 octets)	OctetString6
MacAddressB ^a	MAC address of the source node (6 octets) as seen over adapter B, as advertised by the PRP_Supervision frame	OctetString6
CntReceivedA	Number of frames received from that source over LAN_A	Unsigned32
CntReceivedB	Number of frames received from that source over LAN_B	Unsigned32
CntKeptFramesA	Number of frames that were kept because they were out of the drop window on LAN_A	Unsigned32
CntKeptFramesB	Number of frames that were kept because they were out of the drop window on LAN_B	Unsigned32
CntErrOutOfSequenceA	Number of frames that were out of sequence on LAN_A	Unsigned32
CntErrOutOfSequenceB	Number of frames that were out of sequence on LAN_B	Unsigned32
CntErrWrongLanA	Number of frames that were received with the wrong LAN identifier on LAN_A	Unsigned32
CntErrWrongLanB	Number of frames that were received with the wrong LAN identifier on LAN_B	Unsigned32
TimeLastSeenA	UTC time at which the latest frame was received over LAN_A	UTCTime
TimeLastSeenB	UTC time at which the latest frame was received over LAN_B	UTCTime
SanA	True if the remote device is most probably a SAN accessible over adapter A	Boolean1
SanB	True if the remote device is most probably a SAN accessible over adapter B	Boolean1
SendSeq	Sequence number used to communicate with that remote device	Unsigned16
^a MacAddressB is not a key attribute.		

4.3.3 PRP write

This service shall be used to write values to the LRE of a DANP to control the PRP. Table 5 shows the parameters of this service.

Table 5 – PRP write

Parameter name	Req	Ind	Rsp	Cnf
Argument	M	M(=)		
NodeName	M	M(=)		
ManufacturerName	M	M(=)		
VersionName	M	M(=)		
MacAddressA	M	M(=)		
MacAddressB	M	M(=)		
AdapterActiveA	M	M(=)		
AdapterActiveB	M	M(=)		
DuplicateDiscard	M	M(=)		
TransparentReception	M	M(=)		
SwitchingEndNode	M	M(=)		
NodesTableClear	M	M(=)		
Supervision Address	M	M(=)		
LifeCheckInterval	U	U(=)		

Parameter name	Req	Ind	Rsp	Cnf
NodeForgetTime	U	U(=)		
DropWindowMax	U	U(=)		
Result (+)			S	S(=)
Status			M	M(=)
Result (-)			S	S(=)
Status			M	M(=)

NOTE For the meaning of Req, Ind, Rsp, Cnf, M, U and S, refer to ISO/IEC 10164-1.

Argument

The argument shall convey the service specific parameters of the service request as defined in 4.3.1.

Result(+)

This parameter indicates that the service request succeeded.

Status

This parameter shall return 0 (no error condition detected)

Result(-)

This parameter indicates that the service request failed

Status

This parameter specifies the error condition (see MIB in Clause 7)

4.3.4 PRP read

This service shall be used to read the current status of the LRE from a DANP. Table 6 shows the parameters of this service.

Table 6 – PRP read

Parameter name	Req	Ind	Rsp	Cnf
Argument	M	M(=)		
none				
Result (+)			S	S(=)
NodeName			M	M(=)
ManufacturerName			M	M(=)
VersionName			M	M(=)
MacAddressA			M	M(=)
MacAddressB			M	M(=)
AdapterActiveA			M	M(=)
AdapterActiveB			M	M(=)
DuplicateDiscard			M	M(=)
TransparentReception			M	M(=)
SwitchingEndNode			M	M(=)
NodesTableClear			M	M(=)
SupervisionAddress			M	M(=)
LifeCheckInterval			M	M(=)
NodeForgetTime			M	M(=)
DropWindowMax			M	M(=)

Parameter name	Req	Ind	Rsp	Cnf
CntTotalSentA			M	M(=)
CntTotalSentB			M	M(=)
CntErrorsA			M	M(=)
CntErrorsB			M	M(=)
CntNodes			M	M(=)
NodesTable			M	M(=)
Result (-)			S	S(=)
Status			M	M(=)

NOTE For the meaning of Req, Ind, Rsp, Cnf, M, U and S, refer to ISO/IEC 10164-1.

Argument

The argument shall convey the service specific parameters of the service request as defined in 4.3.1.

Result(+)

This parameter indicates that the service request succeeded.

Result(-)

This parameter indicates that the service request failed.

Status

This parameter specifies the error condition (see MIB in Clause 7).

5 High-availability Seamless Redundancy (HSR)**5.1 HSR objectives**

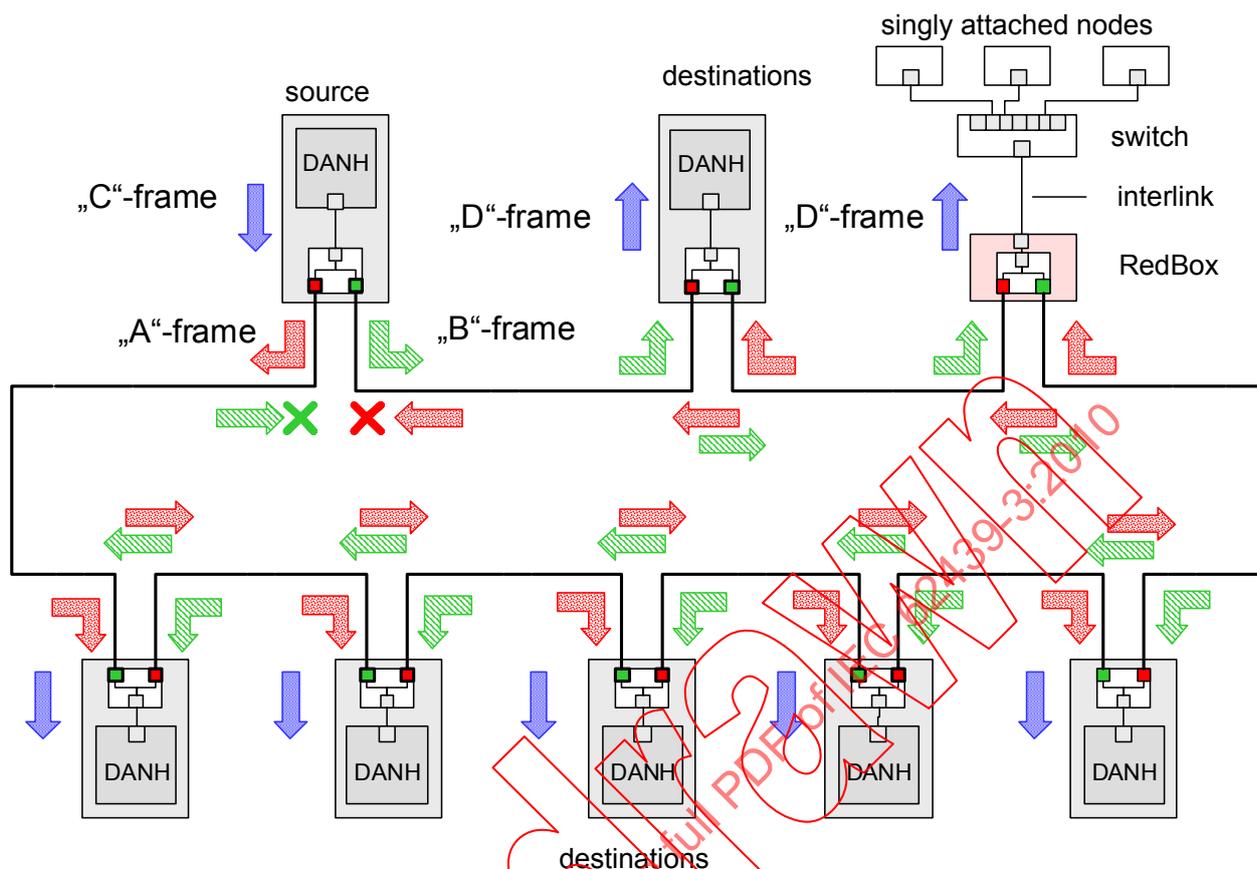
Clause 5 describes the application of the PRP principles (Clause 4) to implement a High-availability Seamless Redundancy (HSR), retaining the PRP property of zero recovery time, applicable to any topology, in particular rings and rings of rings.

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) or IEC 62439-6 (DRP), the available network bandwidth for network traffic is roughly halved. Nodes within the ring are restricted to be HSR-capable switching end nodes. General-purpose nodes (SANs) cannot be attached directly to the ring, but need attachment through a RedBox (redundancy box).

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 switching nodes, each having two ring ports, interconnected by full-duplex links, as shown in the example of Figure 13 (multicast) and Figure 14 (unicast) for a ring topology.



IEC 368/10

Key

red, dotted arrows
green, cross-hatched arrows
blue arrows
cross

“A” frames
“B” frames
non-HSR frames exchanged between ring and host
frame is removed from the ring by the next node

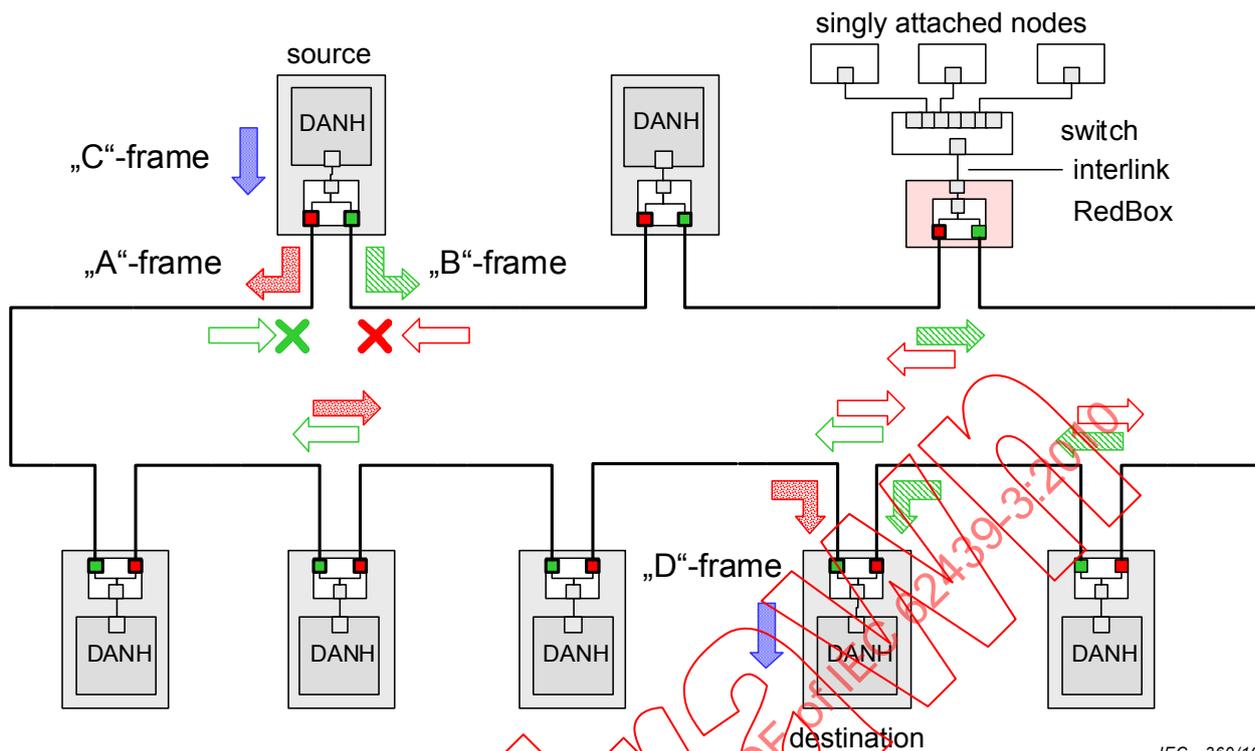
Figure 13 – HSR example of ring configuration for multicast traffic

A source DANH sends a frame passed from its upper layers (“C” frame), inserts an HSR tag to identify frame duplicates and sends a 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.

In particular, the node will not forward a frame that it injected into the ring.



IEC 369/10

Key

- red dotted arrows
- green cross-hatched arrows
- blue arrows
- cross

- „A“ frames
- „B“ frames
- non-HSR frames exchanged between ring and host
- frame is removed from the ring by the next node

Figure 14 – 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.

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 The 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 8802-2 (IEEE 802.2) definition), each node could introduce 125 µs of delay at 100 Mbit/s. With 50 nodes, the time skew may exceed 6 ms.

5.2.2 DANH node structure

Figure 15 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 switching. 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.

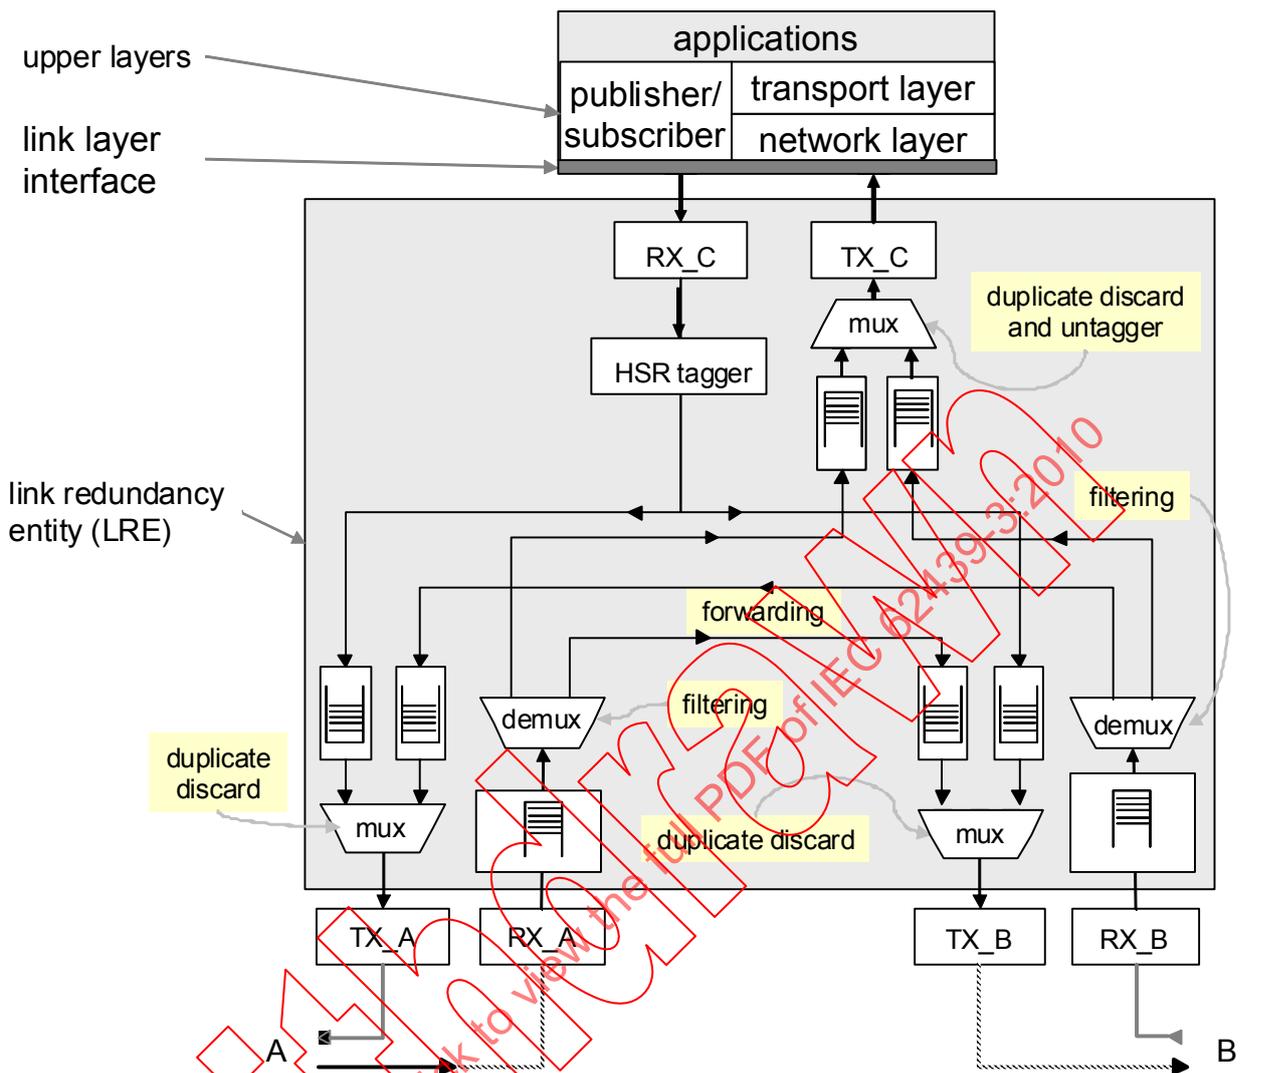


Figure 15 –HSR structure of a DANH

IEC 370/10

5.2.3 Topology

5.2.3.1 Attachment of singly attached nodes

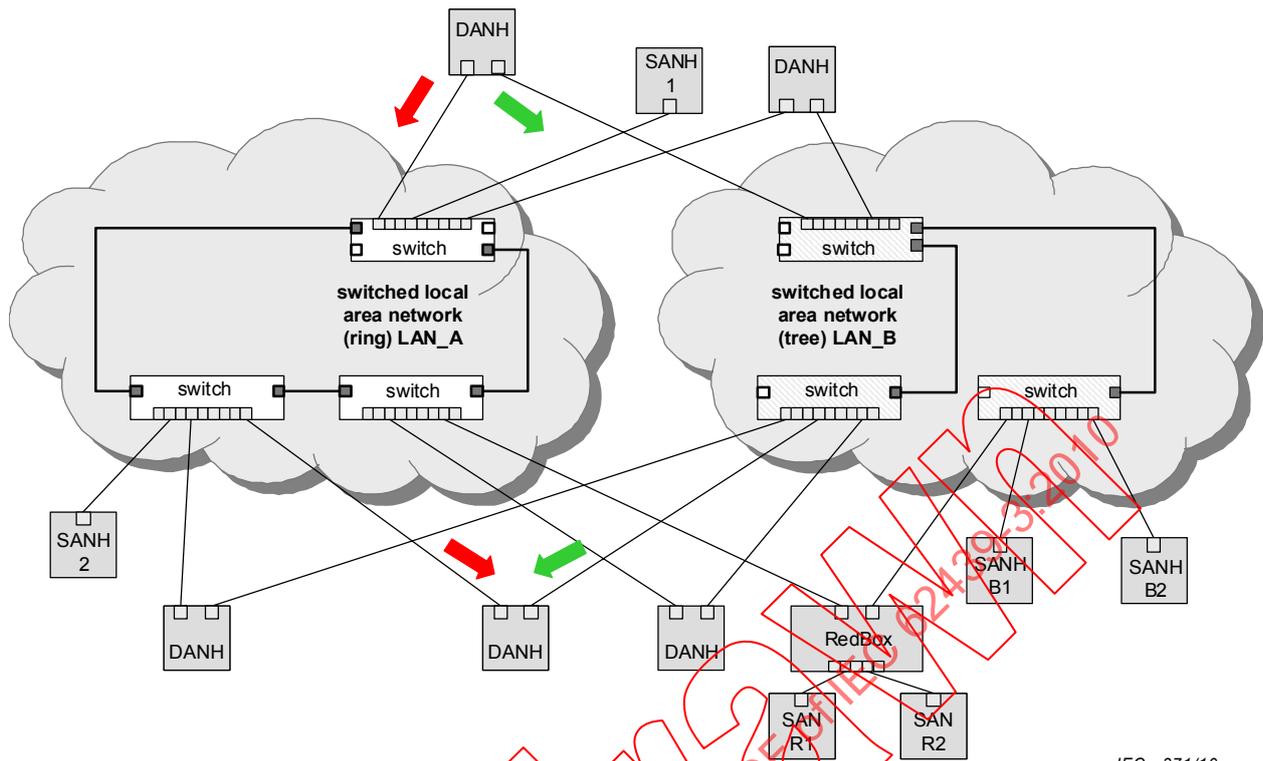
Singly attached nodes (SAN), for instance maintenance laptops or printers cannot be inserted directly into the ring since they have only one port and cannot interpret the HSR tag in the frames. SANs communicate with ring devices through a RedBox (redundancy box) that acts as a proxy for the SANs attached to it, as shown in Figure 13 and Figure 14. The RedBox is detailed in 5.2.4.

Connecting non-HSR nodes to ring ports, breaking the ring, is not covered by this document. Non-HSR traffic within the ring is not permitted.

5.2.3.2 Use of HSR with separate LANs

HSR nodes can be connected in the same way as PRP nodes. In this case, the HSR nodes do not forward frames from port to port (HSR_PRP mode).

However, SANs cannot be attached directly to such a duplicated network unless they are able to interpret the HSR tag. See Figure 16 below.



IEC 371/10

Figure 16 – 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 17 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.

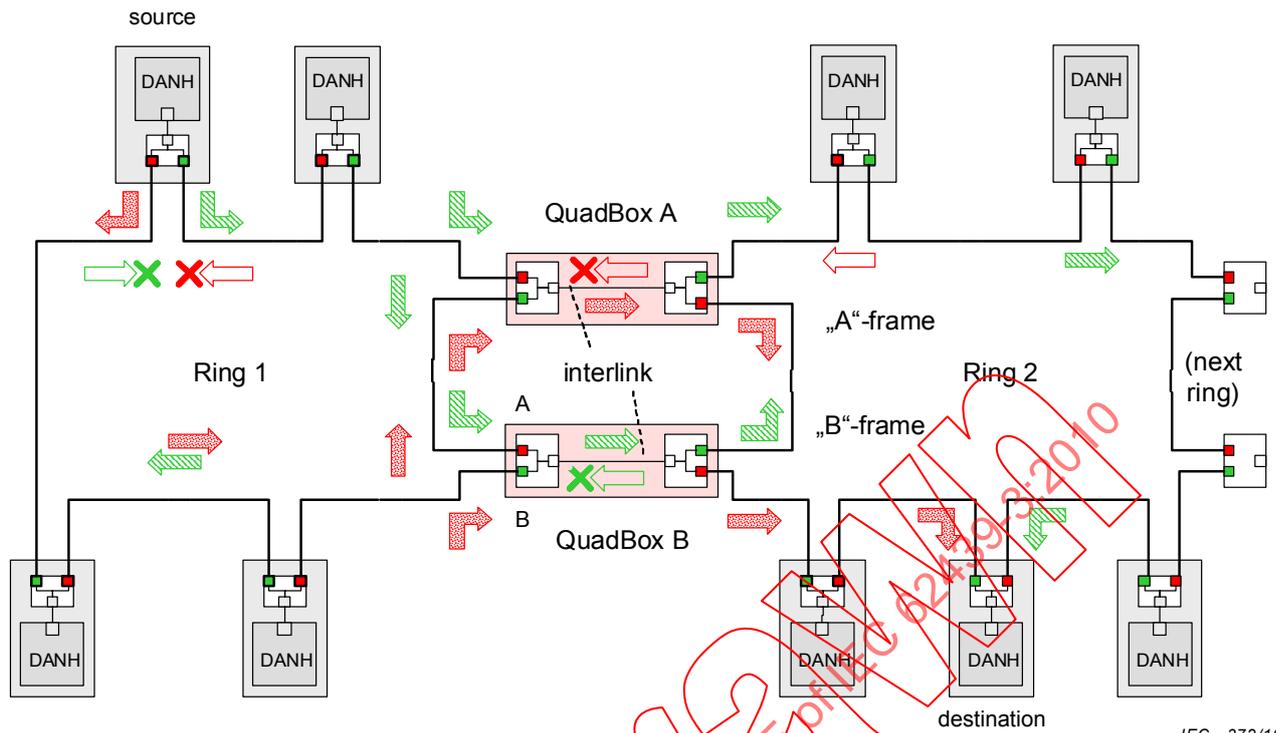


Figure 17 – HSR example of peer coupling of two rings

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 other two copies.

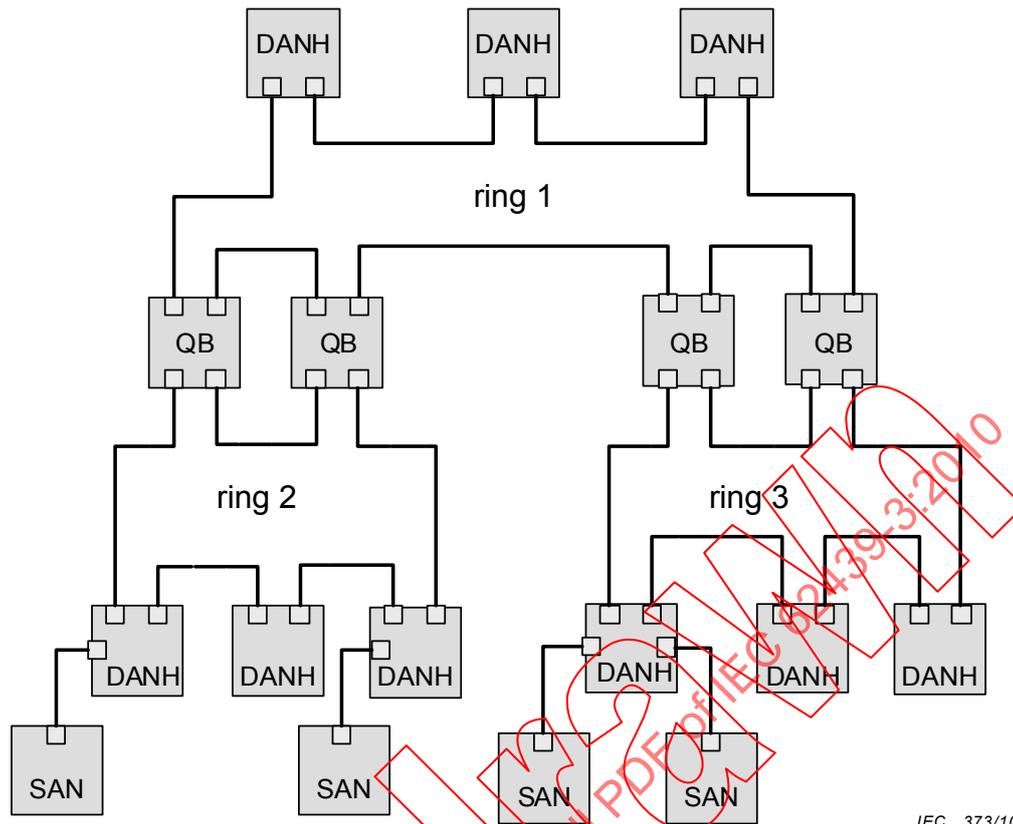
This does not cause four frames to circulate on the second ring, since, when a copy from a first QuadBox reaches the second QuadBox on the same second ring, the second QuadBox will not forward it if it already sent a copy that came from its interlink. Conversely, if the second QuadBox did not yet receive 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 18 shows.



IEC 373/10

Figure 18 – 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

A HSR may be coupled to a PRP network through two RedBoxes, one for each LAN, as Figure 19 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 communication associations to persist through the translation 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.

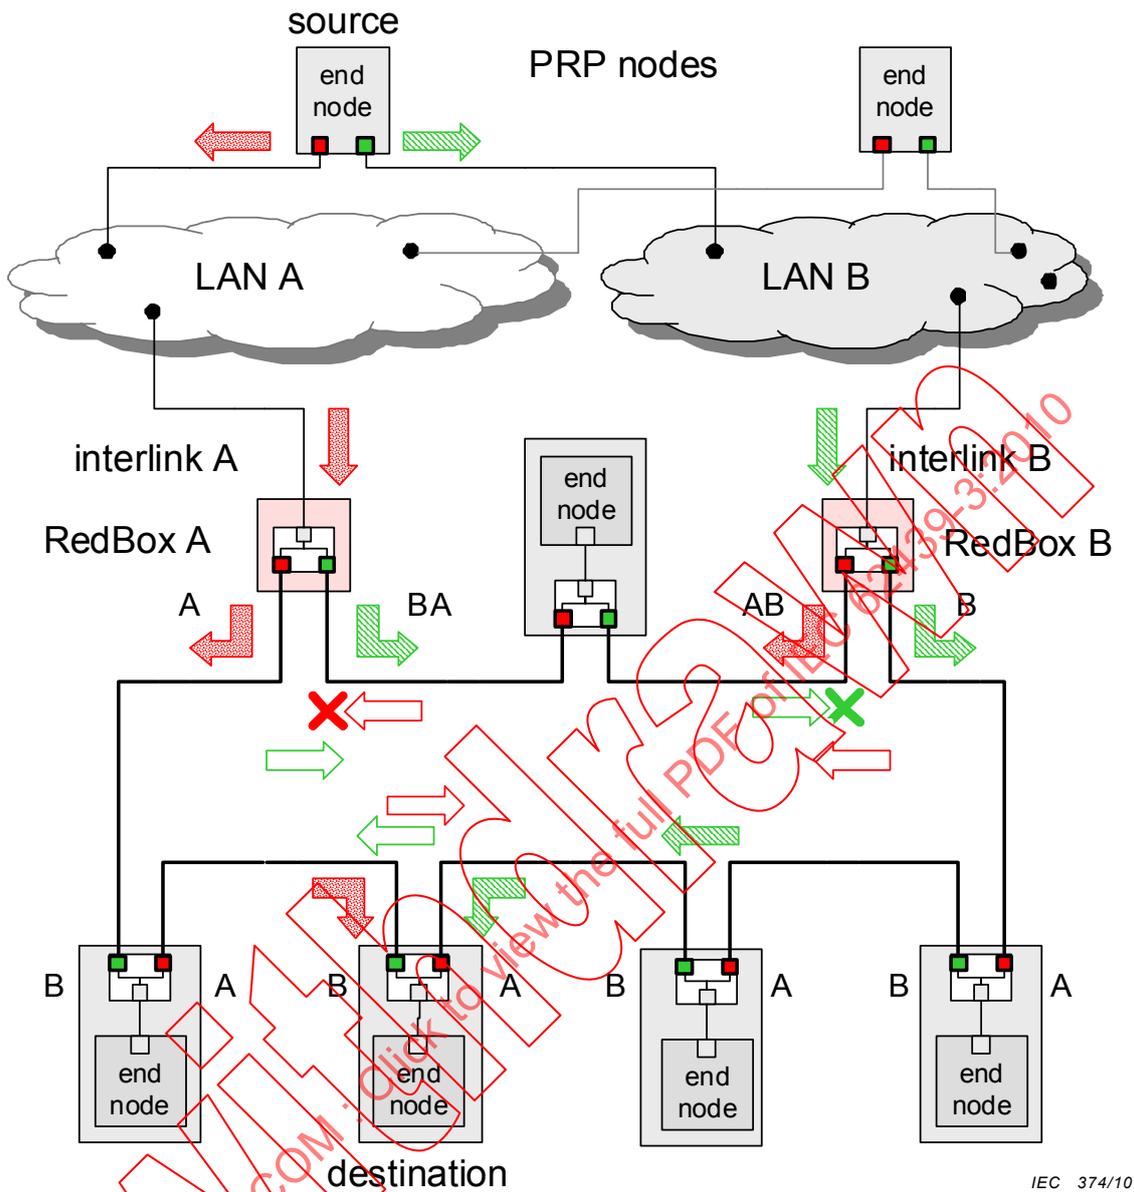


Figure 19 – HSR example of coupling two redundant PRP LANs to a ring

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

In the example of Figure 19, 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 19) are removed by the RedBox that inserted them into the ring, if they originated from outside the ring.

Figure 20 shows the same coupling when the source is within the ring.

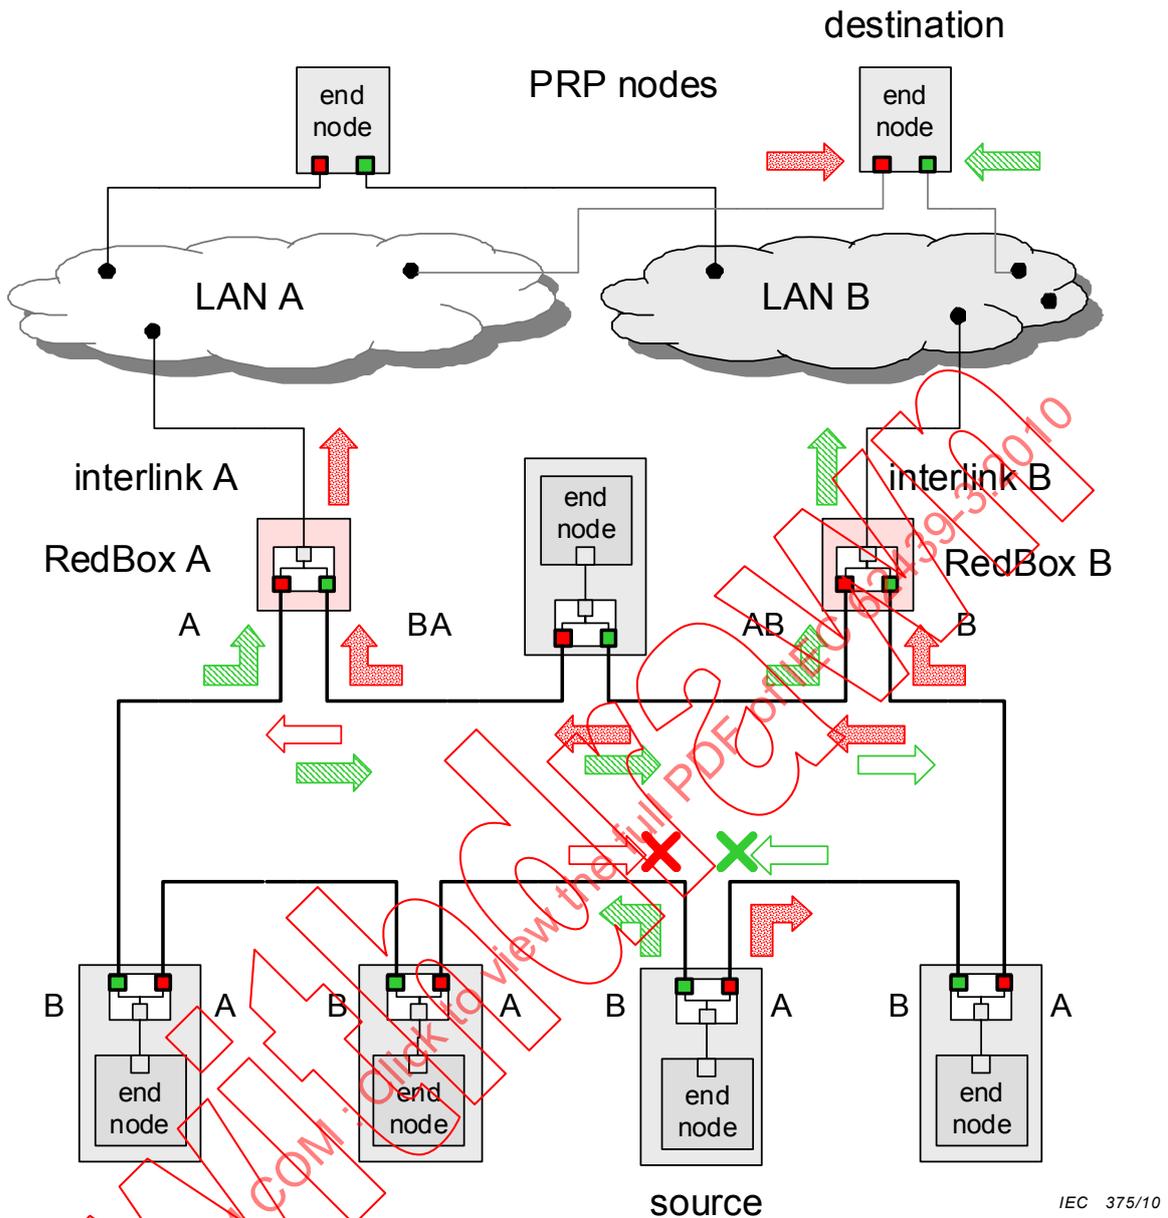
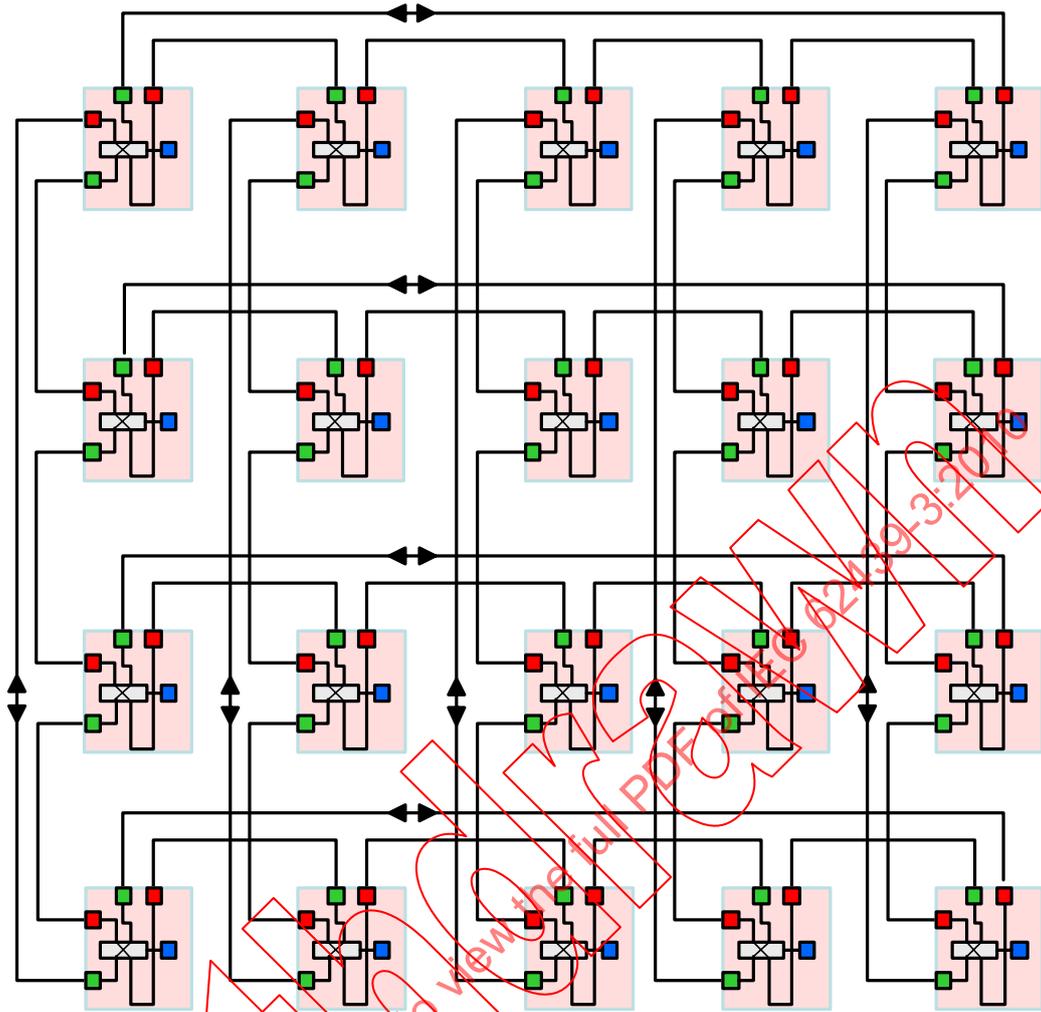


Figure 20 – HSR example of coupling from a ring node to redundant PRP LANs

It is necessary to configure the RedBox as a connexion to SAN or to PRP since the RedBox must insert the PRP trailer. However, letting the RedBox operate always in PRP mode does not harm, since the PRP trailer is invisible to the SANs.

5.2.3.6 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 21 shows for 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.

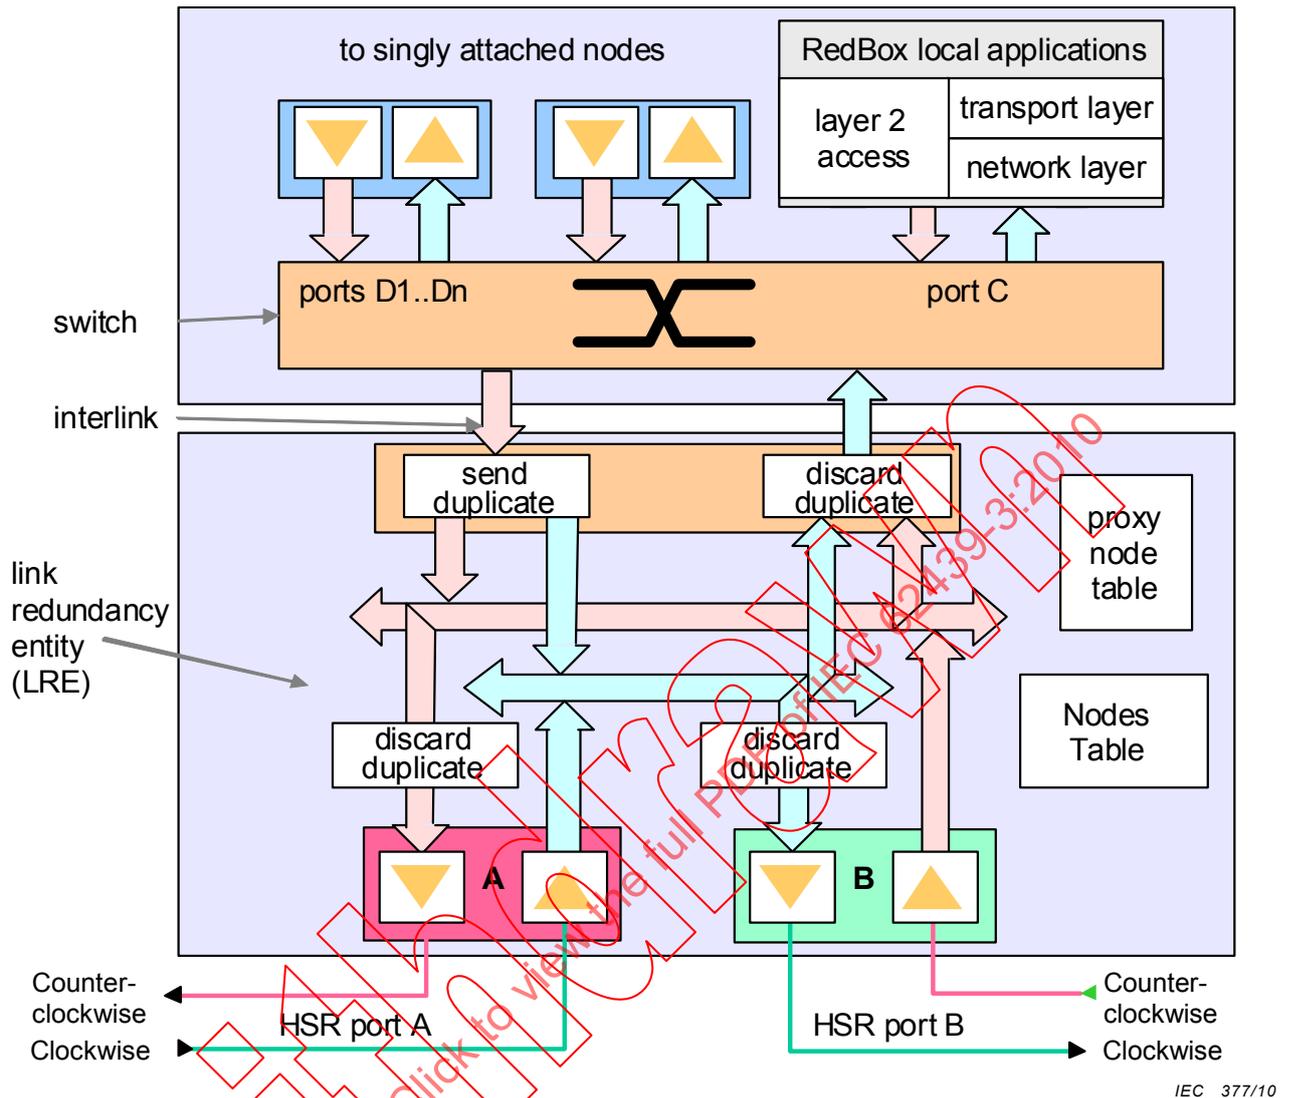


IEC 376/10

Figure 21 – HSR example of meshed topology

5.2.4 RedBox structure

Figure 22 shows the general structure of a RedBox.



IEC 377/10

Figure 22 – HSR structure of a RedBox

The RedBox has a LRE that performs the duties of the HSR protocol.

A RedBox may be operated in three modes: as a SAN, a PRP or as an HSR connection. Depending on the mode of operation, the frame handling at the interlink interface of the RedBox differs.

The RedBox receives the frames to be sent from its own upper layers or from other nodes over its interlink.

The RedBox registers the presence of the source node in its proxy node table, and, if the node does not yet exist, it creates an entry for that node with a sequence number of 0.

If the frame has an HSR tag, the Redbox does not modify it.

If the frame has a PRP trailer, the RedBox reuses the PRP sequence number of the RCT for the HSR tag.

If the frame has no PRP trailer, the RedBox uses the sequence number of its proxy node table and increments it.

The RedBox forwards the frames received from one port to the other, unless the frame was sent already or was sent on behalf of one of the nodes registered in the proxy node table.

The RedBox receives frames addressed to its own upper protocols or to one of the devices that it represents.

If the destination node has been registered as an HSR node, the RedBox forwards the frame unchanged.

If the destination node has been registered as PRP, the RedBox removes the HSR tag and inserts the PRP tag, reusing the sequence number of the HSR tag for the RCT.

If the destination node is neither a PRP nor an HSR node, the RedBox removes the HSR tag.

The switch in Figure 22 may 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 Host sequence number

An HSR node shall maintain a sequence number on behalf of its host 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 in the LRE. Without knowledge of the host MAC address, the LRE forwards all HSR traffic to the host, treating unicast frames as multicast.

5.3.2 DANH receiving from its link layer interface

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

- If this frame is HSR or the destination node has been registered as non-*HSR*
 - Do not modify the frame;
- else (non-*HSR* frame and destination node not registered as non-*HSR*)
 - 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 not cause circulating frames since such frames will not be forwarded by the adjacent HSR node. This mechanism is intended to allow 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:

Register the source in its node table as non-HSR node;

Enqueue the unchanged frame for passing to its link layer interface.

(the frame is not forwarded)

Else (HSR-tagged frame):

Register the source in its node table as HSR node;

If this node is the (unicast or multicast) destination:

If this is the first occurrence of the frame over the link layer interface:

Register the occurrence of that frame;

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):

Register the occurrence of that frame;

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):

If this is the first occurrence of the frame over the second port:

Register the occurrence of that frame;

Enqueue the unmodified frame for sending over the second port.

Else (this is not the first occurrence of the frame over the second port):

Register the occurrence of that frame;

Discard the frame.

Else (If this node is the only (unicast) destination):

Discard the frame.

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

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 shall not send back a frame over the port which received it.

A node that detects on the base of the signal quality supervision 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 27.3.1.4.2 of ISO/IEC 8802-3:2000 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 only allow buffering when the port is reconnected.

If a node receives a supervision frame from a previously connected node indicating reinitialization, it shall purge the buffers from the entries corresponding to that node.

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, there may be a delay of some 6 ms until a frame comes back to its originator, so this possibility must be cared for.

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 For cut-through operation, the node must wait approximately 5 μ s at 100 Mbit/s until the HSR tag has been completely received and the node decided to forward or not. By contrast, store-and-forward takes at least 122 μ s 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

HSR does not specify the clock synchronization method that has to be used.

In case IEC 61588 (IEEE 1588) is used, and due to the fact that clock synchronization frames can arrive on both ports with different delays, it is recommended to use transparent – ordinary clocks (hybrid clocks) together with one-step mode, according to IEC 61588 (IEEE 1588) v2, the P-delay measurement being done directly between adjacent nodes.

NOTE One-step clocks require on-the-fly modification of the clock correction, which is only practical when done in hardware.

5.3.7 Deterministic medium access

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.

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 be configurable for one of three modes:

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

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

NOTE 1 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 2 The interlink can be an internal connection if the RedBox serves as switch 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:

Register the source as an HSR source;

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

Register the occurrence

Discard the frame

Else (If this is the first occurrence of the frame at the link layer interface)

Register the occurrence

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 at each HSR port

Register the occurrence

Discard the frame (already sent over that port)

Else (If this is the first occurrence of the frame at each HSR port)

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

If the source MAC address is not already registered:

Create an entry in the proxy node table;

Register that source as PRP "A" or "B" (depending on the PRP mode of the RedBox)

If the PRP tag does not correspond to the mode of the RedBox "A" or "B"

Register the error;

Discard the frame

Else (If the PRP tag corresponds to the mode of the RedBox "A" or "B")

If the PRP frame was already received

Register the occurrence

Discard the frame

Else (if the PRP frame was not already received)

Register the occurrence

If the RedBox is a destination of the frame

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

Register the occurrence

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 path identifier to build the HSR tag
Enqueue the unmodified frame into each HSR port

Else (if the frame carries neither a HSR tag nor a PRP RCT)

If the source MAC address is not already registered:

Create an entry in the proxy node table with a sequence number of 0;

Register that source as SAN

Else (If the source is already registered)

Register the presence of that source;

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)

If this is not the first occurrence of the frame at each HSR port

Register the occurrence

Discard the frame (already sent over that port)

Else (If this is the first occurrence of the frame at each HSR port)

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 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.

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 proxy node table. This condition is enabled by default and can be disabled for debugging purposes.

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:

Register the source as a non-HSR source;

Enqueue the frame for passing to the link layer interface of the RedBox;

(do not forward it)

Else (frame is HSR-tagged):

If this is the first occurrence of the frame in direction of the second HSR port

Register the occurrence of the frame;

Enqueue the frame to the second port;

Else (If this is not the first occurrence of the frame in direction of the second HSR port)

Register the occurrence of the frame;

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:

Register the occurrence of the frame;

If the RedBox is in SAN mode:

Remove the HSR tag;

Else if the RedBox is in PRP mode:

Remove the HSR tag and append the PRP RCT "A" or "B" reusing the HSR sequence number.

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):

Register the occurrence;

Discard the frame;

NOTE 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 proxy node table handling

A RedBox shall hold a proxy node table containing an entry for each represented association, which shall support at least ProxyTableMaxEntries entries.

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

NOTE The actual size of the proxy node table is indicated in the PICS.

5.4.6 RedBox CoS

Same as 5.3.5, except that 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.

5.4.7 RedBox clock synchronization

Same as 5.3.6.

5.4.8 RedBox medium access

Same as 5.3.7, except that the RedBox must 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 Association definition

For the purpose of duplicate discard, a frame shall be identified by:

- its source MAC address;
- a sequence number;
- the VLAN identifier.

NOTE It is possible to use other fields of the frame such as the checksum to aid in duplicate detection.

5.7 Frame format for HSR

5.7.1 HSR-tagged frame format

Frames to be treated HSR are identified uniquely by their source MAC address, destination MAC address and HSR tag.