



BSI Standards Publication

Industrial communication networks — High availability automation networks

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

NO COPYING WITHOUT BSI PERMISSION EXCEPT AS PERMITTED BY COPYRIGHT LAW

raising standards worldwide™

National foreword

This British Standard is the UK implementation of EN 62439-3:2012. It is identical to IEC 62439-3:2012. It supersedes BS EN 62439-3:2010 which is withdrawn.

The UK participation in its preparation was entrusted to Technical Committee AMT/7, Industrial communications: process measurement and control, including fieldbus.

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

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

© The British Standards Institution 2012

Published by BSI Standards Limited 2012

ISBN 978 0 580 73809 8

ICS 25.040.01; 35.040

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

This British Standard was published under the authority of the Standards Policy and Strategy Committee on 31 October 2012.

Amendments issued since publication

Amd. No.	Date	Text affected
----------	------	---------------

EUROPEAN STANDARD
NORME EUROPÉENNE
EUROPÄISCHE NORM

EN 62439-3

September 2012

ICS 25.040; 35.040

Supersedes EN 62439-3:2010

English version

**Industrial communication networks -
High availability automation networks -
Part 3: Parallel Redundancy Protocol (PRP) and High-availability
Seamless Redundancy (HSR)
(IEC 62439-3:2012)**

Réseaux industriels de communication -
Réseaux d'automatisme à haute
disponibilité -
Partie 3 : Protocole de redondance
parallèle (PRP) et redondance
transparente de haute disponibilité (HSR)
(CEI 62439-3:2012)

Industrielle Kommunikationsnetze -
Hochverfügbare Automatisierungsnetze -
Teil 3: Parallelredundanz-Protokoll (PRP)
und nahtloser Hochverfügbarkeits-Ring
(HSR)
(IEC 62439-3:2012)

This European Standard was approved by CENELEC on 2012-08-09. CENELEC members are bound to comply with the CEN/CENELEC Internal Regulations which stipulate the conditions for giving this European Standard the status of a national standard without any alteration.

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

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

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

CENELEC

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

Management Centre: Avenue Marnix 17, B - 1000 Brussels

© 2012 CENELEC - All rights of exploitation in any form and by any means reserved worldwide for CENELEC members.

Ref. No. EN 62439-3:2012 E

 **Documents**
Document Sharing Hub

Foreword

The text of document 65C/687/FDIS, future edition 2 of IEC 62439-3, prepared by SC 65C, "Industrial networks", of IEC TC 65, "Industrial-process measurement, control and automation" was submitted to the IEC-CENELEC parallel vote and approved by CENELEC as EN 62439-3:2012.

The following dates are fixed:

- latest date by which the document has to be implemented at national level by publication of an identical national standard or by endorsement (dop) 2013-05-09
- latest date by which the national standards conflicting with the document have to be withdrawn (dow) 2015-08-09

This document supersedes EN 62439-3:2010.

EN 62439-3:2012 includes the following significant technical changes with respect to EN 62439-3:2010:

- specification of the interconnection of PRP and HSR networks;
- introduction of a suffix for PRP frames;
- clarification and modification of specifications to ensure interoperability;
- slackening of the specifications to allow different implementations;
- consideration of clock synchronization according to IEC 61588;
- introduction of test modes to simplify testing and maintenance.

This standard is to be used in conjunction with EN 62439-1:2010.

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

Endorsement notice

The text of the International Standard IEC 62439-3:2012 was approved by CENELEC as a European Standard without any modification.

In the official version, for Bibliography, the following note has to be added for the standard indicated:

IEC 61580 series NOTE Harmonized in EN 61580 series (not modified).

Annex ZA
(normative)**Normative references to international publications
with their corresponding European publications**

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

NOTE When an international publication has been modified by common modifications, indicated by (mod), the relevant EN/HD applies.

<u>Publication</u>	<u>Year</u>	<u>Title</u>	<u>EN/HD</u>	<u>Year</u>
IEC 60050-191	-	International Electrotechnical Vocabulary (IEV) - Chapter 191: Dependability and quality of service	-	-
IEC 61588	-	Precision clock synchronization protocol for networked measurement and control systems	-	-
IEC 62439-1	-	Industrial communication networks - High availability automation networks - Part 1: General concepts and calculation methods	EN 62439-1	-
IEC 62439-2	-	Industrial communication networks - High availability automation networks - Part 2: Media Redundancy Protocol (MRP)	EN 62439-2	-
IEC 62439-6	-	Industrial communication networks - High availability automation networks - Part 6: Distributed Redundancy Protocol (DRP)	EN 62439-6	-
IEC 62439-7	-	Industrial communication networks - High availability automation networks - Part 7: Ring-based Redundancy Protocol (RRP)	EN 62439-7	-
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 and Metropolitan Area Networks - Media Access Control (MAC) Bridges	-	-
IEEE 802.1Q	2011	IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks	-	-

CONTENTS

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

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.




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.




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

IEC 360/10

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

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

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.