INTERNATIONAL STANDARD

ISO 9542

Second edition
19xx-yy-zz

INTERNATIONAL ORGANIZATION FOR STANDARDIZATION
ORGANISATION INTERNATIONALE DE NORMALISATION

Information processing systems - Telecommunications and information exchange between systems - End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)

Systèmes de traitement de l'information — Téléinformatique — Protocole de routage d'un système d'etrémité à un système intermédaire à utiliser conjointement avec le protocole fournissant le service de réseau en mode sans connexion (ISO 8473)

Editor's Note:

This document is a working draft incorporating all current defect report corrections to the base standard.

It is **not** and approved ISO Standard

Reference number
ISO 9542: 19xx (E)

Foreword

ISO (the International Organization for Standardization) is a worldwide federation of Inational standards bodies (ISO member bodies). The word of preparing International Standards is normally carried out through ISO technical committees. Each member body interested in a subject for which a technical committee has been established has the right to be represented on that committee. International organizations, governmental and non-governmental, in liaison with ISO, also take part in the work. ISO collaborates closely with the International Electrotechnical Commission (IEC) on all matters of electrotechnical standardization.

Draft International Standards adpted by the technical committees are circulated to the member bodies for approval before their acceptance as Iternational Standards by ISO Council. they are approved in accordance with ISO procedures requiring at least 75% approval by the member bodies voting.

Internation Standard ISO 1234 was prepared by Technical Committee ISO//IEC JTC1, Information Processing Systems.

Users should note that all International Standards undergo revision from time to time and that any reference made herein to any other International Standard implies its latest edition, unless otherwise stated.

Contents

Figures

Tables

Introduction

This International Standard is one of a set of International Standards produced to facilitate the interconnection of open systems. The set of standards covers the services and protocols required to achieve such interconnection.

This International Standard is positioned with respect to other related standards by the layers defined in ISO 7498 and by the structure defined in ISO 8648. In particular, it is a protocol of the Network Layer. This International Standard permits End Systems and Intermediate Systems to exchange configuration and routeing information to facilitate the operation of the routeing and relaying functions of the Network Layer.

The aspects of Network Layer routeing that are concerned with communication between End Systems and Intermediate Systems on the same subnetwork are to a great extent separable from the aspects that are concerned with communication among the Intermediate Systems that connect multiple subnetworks. This protocol addresses only the former aspects. It will be significantly enhanced by the cooperative operation of an additional protocol that provides for the exchange of routeing information among Intermediate Systems, but is useful whether or not such an additional protocol is available.

This International Standard is designed to operate in close conjunction with ISO 8473 and its addenda.

This International Standard provides solutions for the following practical problems.

1) How do End Systems discover the existence and reachability of Intermediate Systems that can route NPDUs to destinations on subnetworks other than the one(s) to which the End System is directly connected?

2) How do End Systems discover the existence and reachability of other End Systems on the same subnetwork (when direct examination of the destination NSAP address does not provide information about the destination subnetwork address)?

3) How do Intermediate Systems discover the existence and reachability of End Systems on each of the subnetworks to which they are directly connected?

4) How do End Systems decide which Intermediate System to use to forward NPDUs to a particular destination when more than one Intermediate System is accessible?

The protocol assumes that:

1) routeing to a specified subnetwork point of attachment address (SNPA) on the same subnetwork is carried out satisfactorily by the subnetwork itself, but

2) the subnetwork is not, however, capable of routeing on a global basis using the NSAP address alone to achieve communication with a requested destination.[1]

In addition, certain protocol functions assume that:

3) the subnetwork supports broadcast, multicast, or other forms of multi-destination addressing for *n*-way transmission.

1Consequently, it is not possible to use Application Layer communication to carry out the functions of this International Standard.

The protocol is connectionless, and is designed to:

- minimize the amount of a priori state information needed by End Systems before they can begin to communicate with other End Systems;
- minimize the amount of memory needed to store routeing information in end systems; and
- minimize the computational complexity of End System routeing algorithms.

Information processing systems — Telecommunications and information exchange between systems - End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473)

1 Scope and Field of Application

This International Standard specifies a protocol which is used by Network Layer entities operating ISO 8473 in End Systems and Intermediate Systems (referred to herein as ES and IS respectively) to maintain routeing information. The Protocol herein described relies upon the provision of a connectionless-mode underlying service.[2]

This International Standard specifies:

1) procedures for the transmission of configuration and routeing information between Network entities residing in End Systems and Network entities residing in Intermediate Systems;

2) the encoding of the protocol data units used for the transmission of the configuration and routeing information;

3) procedures for the correct interpretation of protocol control information; and \item the functional requirements for implementations claiming conformance to this International Standard.

The procedures are defined in terms of:

1) the interactions between End System and Intermediate System Network entities through the exchange of protocol data units; and

2) the interactions between a Network entity and an underlying service provider through the exchange of subnetwork service primitives. \end{enumerate}

This International Standard does *not* specify any protocol elements or algorithms for facilitating routeing and relaying among Intermediate Systems. Such functions are intentionally beyond the scope of this International Standard.

2 References

ISO 7498, *Information processing systems — Open systems interconnection — Basic reference model.*

ISO 7498/Add.1, *Information processing systems — Open systems interconnection — Basic reference model.* ADDENDUM 1: *Connectionless-mode transmission.*

ISO 7498/Add.4, *Information processing systems — Open systems interconnection — Basic reference model*. ADDENDUM 4: *OSI Management Framework.*

ISO 8208, *Information processing systems — Data communications — X.25 Packet Level Protocol for Data Terminal Equipment.*

ISO 8348, *Information processing systems — Data communications — Network Service Definition.*

ISO 8348/Add.1, *Information processing systems — Data communications — Network Service Definition.* ADDENDUM 1: *Connectionless-mode Transmission.*

ISO 8348/Add.2, *Information processing systems — Data communications — Network Service Definition*. ADDENDUM 2: *Network Layer Addressing.*

ISO 8473, *Information processing systems — Data communications — Protocol for providing the connectionless-mode Network Service.*

ISO 8648, *Information processing systems — Open Systems Interconnection — Internal organization of the Network layer.*

ISO 8802, *Information processing systems — Data communications – Local Area Networks.*

CCITT X.25, *Interface Between Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) for Terminals Operating in the Packet Mode and Connected to Public Data Networks by Dedicated Circuit, 1985.*

ISO 10039:1990, *IPS-T&IEBS — MAC Service Definition.*

2See Clause 8 of ISO 8473 for the mechanisms necessary to realize this service on subnetworks based on ISO 8208 and ISO 8802.

3 Definitions

3.1 Reference Model Definitions

ISO 9542 makes use of the following terms defined in ISO 7498.

1) Network layer
2) Network service access point
3) Network service access point address
4) Network entity
5) routeing
6) Network protocol
7) Network relay
8) Network protocol data unit

3.2 Network Layer Architecture Definitions

ISO 9542 makes use of the following terms defined in ISO 8648.

1) Subnetwork
2) End System
3) Intermediate System
4) Subnetwork Service
5) Subnetwork Dependent Convergence Function

3.3 Network Layer Addressing Definitions

ISO 9542 makes use of the following terms defined in ISO 8348/Add.2.

1) Subnetwork address
2) Subnetwork point of attachment
3) Network Protocol Address Information
4) Network Entity Title

3.4 Local Area Network Definitions

ISO 9542 makes use of the following terms defined in ISO 8802.

1) multicast address
2) broadcast medium

3.5 Additional Definitions

For the purposes of this International Standard, the following definitions apply.

3.5.1 Configuration: The collection of End and Intermediate Systems attached to a single subnetwork, defined in terms of the system types, NSAP addresses present, Network Entities present, and the correspondence between systems and SNPA addresses

4 Symbols and Abbreviations

4.1 Data Units

PDU	Protocol Data Unit
SNSDU	Subnetwork Service Data Unit
NPDU	Network Protocol Data Unit
SNPDU	Subnetwork Protocol Data Unit

4.2 Protocol Data Units

ESH PDU	End System Hello Protocol Data Unit
ISH PDU	Intermediate System Hello Protocol Data Unit
RD PDU	Redirect Protocol Data Unit

4.3 Protocol Data Unit Fields

NPID	Network Layer Protocol Identifier
LI	Length Indicator
V/P	Version/Protocol Identifier Extension
TP	Type
CS	Checksum
NETL	Network Entity Title Length Indicator
NET	Network Entity Title
DAL	Destination Address Length Indicator
DA	Destination Address
SAL	Source Address Length Indicator
SA	Source Address
BSNPAL	SN Address Length Indicator of better route to destination
BSNPA	SN Address of better route to destination
HT	Holding Time

4.4 Parameters

CT	Configuration Timer
RT	Redirect Timer
ESCT	Suggested End System Configuration Timer

4.5 Addresses

NSAP	Network Service Access Point
SNPA	Subnetwork Point of Attachment

NPAI Network Protocol Address Information

4.6 Miscellaneous

ES End system
IS Intermediate system
LAN Local area network
PICS Protocol Implementation Conformance Statement
QoS Quality of service
SN Subnetwork

5 Overview of the Protocol

5.1 Information Provided by the Protocol

This International Standard provides two types of information to Network entities which support its operation:

- **Configuration information**, and
- **Route redirection information**

Configuration information permits End Systems to discover the existence and reachability of Intermediate Systems and permits Intermediate Systems to discover the existence and reachability of End Systems. This information allows ESs and ISs attached to the same subnetwork to dynamically discover each other's existence and availability, thus eliminating the need for manual intervention at ESs and ISs to establish the identity of Network entities that can be used to route NPDUs.

Configuration information also permits End Systems to obtain information about each other in the absence of an available Intermediate System.

NOTE 1 The term "configuration information" is not intended in the broad sense of configuration as used in the context of OSI system management. Rather, only the functions specifically defined herein are intended.

Route redirection information allows Intermediate Systems to inform End Systems of (potentially) better paths to use when forwarding NPDUs to a particular destination. A better path could either be another IS on the same subnetwork as the ES, or the destination ES itself, if it is on the same subnetwork as the source ES. Allowing the ISs to inform the ESs of routes minimizes the complexity of routeing decisions in End Systems and improves performance because the ESs may make use of the better IS or local subnetwork access for subsequent transmissions.

5.2 Addressing

The **Source Address** and **Destination Address** parameters referred to in this International Standard are OSI Network Service Access Point Addresses. The syntax and semantics of an OSI Network Service Access Point Address are described in ISO 8348/Add.2.

5.3 Underlying Service Assumed by the Protocol

The underlying service required to support this International Standard is defined by the primitives in table 1.

Table 1 - Service Primitives for Underlying Service

SN_UNITDATA	.Request .Indication	SN_Destination_Address, SN_Source_Address SN_Quality_of_Service SN_Userdata

NOTE 2 These service primitives are used to describe the abstract interface which exists between the protocol machine and an underlying real subnetwork or a Subnetwork Dependent Convergence Function which operates over a real subnetwork or real data link to provide the required underlying service.[3]

5.3.1 Subnetwork Addresses

The source and destination addresses specify the points of attachment to a public or private subnetwork(s) involved in the transmission (known as Subnetwork Points of Attachment, or SNPAs). Subnetwork addresses are defined in the service definition of each individual subnetwork.

This International Standard is designed to take advantage of subnetworks which support *broadcast*, *multicast*, or other forms of multi-destination addressing for *n*-way transmission. It is assumed that the SN_Destination_Address parameter may take on one of the following multi-destination addresses in addition to a normal single destination address:

- All End System Network entities
- All Intermediate System Network entities

Where a real subnetwork does not inherently support broadcast or other forms of transmission to multi-destination addresses, a convergence function may be used to provide n-way transmission to these multi-destination addresses.

When the SN_Destination_Address on the SN_UNITDATA.Request is a multi-destination address, the SN_Destination_Address parameter in the corresponding SN_UNITDATA.Indication shall be the same multi-destination address.

The syntax and semantics of subnetwork addresses, except for the properties described above, are not defined in this International Standard.

3See Clause 8 of ISO 8473 for the mechanisms necessary to realize this service on subnetworks based on ISO 8208 and ISO 8802.

5.3.2 Subnetwork User Data

The SN_Userdata is an ordered multiple of octets, and is transferred transparently between the specified subnetwork points of attachment.

The underlying service is required to support a service data unit size of at least that required to operate ISO 8473.

5.4 Subnetwork Types

In order to evaluate the applicability of this International Standard in particular configurations of End Systems, Intermediate Systems and subnetworks, three generic types of subnetwork are identified. These are:

1) the **point-to-point** subnetwork,

2) the **broadcast** subnetwork, and

3) the **general topology** subnetwork

These subnetwork types are discussed in the following clauses.

5.4.1 Point-to-Point Subnetworks

A *point-to-point* subnetwork supports exactly two systems. The two systems may be either two End Systems, or an End System and a single Intermediate System. A single point-to-point data link connecting two Network entities is an example of a point-to-point subnetwork.

5.4.1.1 Configuration information on a point-to-point Subnetwork

On a point-to-point subnetwork the configuration information of this International Standard informs the communicating Network entities of the following:

1) whether the topology consists only of two End Systems, or

2) one of the two systems is an Intermediate System.

NOTE 3 On a point-to-point subnetwork, if both systems are Intermediate Systems, then this International Standard is inapplicable to the situation, since an IS-to-IS protocol should be employed instead. However, there is no reason why the configuration information could not be employed in an IS-to-IS environment to ascertain the topology and initiate operation of an IS-to-IS protocol. \end{note}

The Intermediate System is informed of the NSAP address(es) supported by the Network entity in the End System. This permits reachability information and routeing metrics concerning these NSAPs to be disseminated to other Intermediate Systems for the purpose of calculating routes to/from this End System.

5.4.1.2 Route redirection information on a point-to-point Subnetwork.

Redirection information is not employed on point-to-point subnetworks because there are never any alternate routes.

5.4.2 Broadcast Subnetworks

A *broadcast* subnetwork supports an arbitrary number of End Systems and Intermediate Systems, and additionally is capable of transmitting a single SNPDU to all or a subset of these systems in response to a single SN_UNITDATA.-Request. An example of a broadcast subnetwork is a LAN (local area network) conforming to ISO 8802-2, type 1 operation.

5.4.2.1 Configuration information on a broadcast Subnetwork

On a broadcast subnetwork the configuration information of this International Standard is employed to inform the communicating Network entities of the following:

1) End Systems are informed of the reachability, Network Entity Title, and SNPA address(es) of each active Intermediate System on the subnetwork.

2) Intermediate Systems are informed of the NSAP addresses supported by each End System and the SNPA address(es) of the ES. Once the Intermediate System obtains this information, reachability information and routeing metrics concerning these NSAPs may be disseminated to other ISs for the purpose of calculating routes to/from each ES on the subnetwork.

3) In the absence of an available Intermediate System, End Systems may query over a broadcast subnetwork to discover whether a particular NSAP is reachable on the subnetwork, and if so, what SNPA address to use to reach that NSAP.

5.4.2.2 Route redirection information on broadcast Subnetworks.

Redirection information may be employed on broadcast subnetworks to permit Intermediate Systems to inform End Systems of superior routes to a destination NSAP. The superior route might be another IS on the same subnetwork as the ES, or it might be the destination ES itself, if it is directly reachable on the same subnetwork as the source ES.

5.4.3 General Topology Subnetworks

A *general topology* subnetwork supports an arbitrary number of End Systems and Intermediate Systems, but does not support a convenient multi-destination connectionless transmission facility as does a broadcast subnetwork. An example of a general topology subnetwork is a subnetwork employing X.25 or ISO 8208.

NOTE 4 The crucial distinguishing characteristic between the broadcast subnetwork and the general topology subnetwork is the "cost" of an *n*-way transmission to a potentially large subset of the systems on the subnetwork. On a general topology subnetwork, the cost is assumed to be close to the cost of sending an individual PDU to *each* SNPA on the subnetwork. Conversely, on a broadcast subnetwork the cost is assumed to be close to the cost of sending a single PDU to *one* SNPA on the subnetwork. Intermediate situations between these extremes are of course

possible. In such cases it would be possible to treat the subnetwork as in either the broadcast or general topology category.

5.4.3.1 Configuration information on a general topology Subnetwork

On a general topology subnetwork the configuration information is generally not employed because the protocol can be very costly in the utilization (and charging for) subnetwork resources.

5.4.3.2 Route redirection information on a general topology Subnetwork.

Redirection information may be employed on general topology subnetworks to permit Intermediate Systems to inform End Systems of superior routes to a destination NSAP. The superior route might be another IS on the same subnetwork as the ES, or it might be the destination ES itself, if it is directly reachable on the same subnetwork as the source ES.

Section two: Specification of the protocol

6 Protocol Functions

This section describes the functions performed as part of the protocol.

Implementations are not required to perform all of the functions: Clause 8.1 specifies which functions are mandatory and which are optional.

6.1 Protocol Timers

Many of the protocol functions are timer based. This means that they are executed upon expiration of a timer rather than upon receipt of a PDU or invocation of a service primitive. The two types of timer employed by the protocol are the Configuration Timer (CT) and the Holding Timer (HT).

NOTE 5 It is recommended that the timer values be implemented with a resolution not worse than one second.

6.1.1 Configuration Timer

The Configuration Timer is a local timer (i.e. maintained independently by each system) which assists in performing the Report Configuration function (see 6.2). The timer determines how often a system reports its availability to the other systems on the same subnetwork. The shorter the Configuration Timer, the more quickly other systems on the subnetwork will become aware when the reporting system becomes available or unavailable. There is a trade off between increased responsiveness and increased use of resources in the subnetwork and in the recipient systems.

6.1.2 Holding Timer

The Holding Timer applies to both configuration information and route redirection information. The value of a Holding Timer is set by the source of the information and transmitted in the Holding Time field of the appropriate PDU. The recipient of the information is expected to retain the information no longer than the Holding Timer. Old configuration or redirection information shall be discarded after the Holding Timer expires to ensure the correct operation of the protocol. A holding time value of zero is permitted. When configuration and/or redirection information with a zero holding time is received, prior information shall be replaced, thus causing the system to set its holding timer to zero and discard the corresponding information.

Further discussion of the rationale for these timers and guidelines for their use may be found in Annex B.

6.2 Report Configuration Function

The Report Configuration Function is used by End Systems and Intermediate Systems to inform each other of their reachability and current subnetwork address(es). This function is invoked every time the local Configuration Timer (CT) expires in an ES or IS. The function may optionally be invoked on other occasions. For example, when one of the system's SNPAs becomes operational, this function may be executed more frequently than on Configuration Timer expiry. This enables other systems to notice the change in configuration quickly.

6.2.1 Report Configuration by End Systems

An End System Network entity constructs and transmits ESH PDUs to inform other systems about the NSAPs it serves. This may be done by constructing one ESH PDU for each NSAP. Alternatively, ESH PDUs may be constructed which convey information about more than one NSAP at a time, up to the limits imposed by the permitted SNSDU size and the maximum header size of the ESH PDU. Each ESH PDU is transmitted by issuing an SN-UNITDATA.Request with the following parameters:

SN_Userdata (SNSDU) ← ESH PDU

SN_Destination_Address ← multi-destination address that indicates "*All Intermediate System Network Entities*".

Where an End System supports more than one SNPA, the information about each NSAP served by the End System shall be transmitted on each SNPA. It is not required that the distribution of NSAPs among ESH PDUs be the same on each SNPA.

NOTE 6 The necessity to inform other systems about individual NSAPs served by the Network entity arises from the lack of a formalized relationship between Network entity titles and NSAP addresses. If this relationship could be constrained to require that all NSAP addresses be assigned as leaf subdomains of a domain represented by the local Network entity's Network entity title, then a single ESH PDU could be transmitted containing the ES's Network entity title. The Network entity title would then imply which NSAPs might be present at that End System.

The Holding Time (HT) field is set to approximately twice the ESs Configuration Timer (CT) parameter. The value shall be large enough so that even if every other ESH PDU is discarded (due to lack of resources), or otherwise lost in the subnetwork, the configuration information will still be maintained. The value should be set small enough so that Intermediate Systems can respond in a timely fashion to End Systems becoming available or unavailable.

NOTE 7 The actual value of the SN_Destination_Address used to mean "*All Intermediate System Network Entities*" is subnetwork dependent and will most likely vary from subnetwork to subnetwork. It is of course desirable on widely-used subnetwork types (such as those based on ISO 8802) that this value, and the value of the "*All End System Network Entities*" multi-destination address, be standardized.

6.2.2 Report Configuration by Intermediate Systems

An Intermediate System constructs a single ISH PDU containing the IS's Network entity title and issues one SN_UNITDATA.Request on each SNPA to which it is attached with the following parameters:

SN_Userdata (SNSDU) ←- ISH PDU

SN_Destination_Address ← multi-destination address that indicates "*All End System Network Entities*".

The Holding Time (HT) field is set to approximately twice the Intermediate System's Configuration Timer (CT) parameter. This variable shall be set to a value large enough so that even if every other ISH PDU is discarded (due to lack of resources), or otherwise lost in the subnetwork, the configuration information will still be maintained. The value should be set small enough so that End Systems will quickly cease to use ISs that have failed, thus preventing "black holes" in the network.

An IS may optionally suggest a value for End Systems on the local subnetwork to use as their Configuration Timers (CT) by including the ESCT option in the transmitted ISH PDU. Setting this option permits an IS to influence the frequency with which ESs transmit ESH PDUs.

NOTE 8 An IS may wish to so influence End Systems in order to trade off the subnetwork resources consumed by the transmission of ESH PDUs against the length of time it is willing to tolerate obsolete configuration information about an ES.

6.3 Record Configuration Function

The Record Configuration function receives ESH or ISH PDUs, extracts the configuration information, and updates the information in the local Network entity's routeing information base.

NOTE 9 If an ES so desires, it may decide to enable the appropriate multi-destination address, thus permitting it to process ESH PDUs multicast by other End Systems. There is potentially some performance improvement to be gained by doing this, at the expense of extra memory, and possibly extra processing cycles in the End System. The ES, by recording other ESs' configuration information, may be able {to route NPDUs directly to ESs on the local subnetwork without first being redirected by an Intermediate System.

Similarly, Intermediate Systems may choose to receive the ISH PDUs of other ISs, allowing this International Standard to be used as the initialization and topology maintenance portion of a full IS-to-IS routeing protocol.

The receiving system is not required to process any option fields in a received ESH or ISH PDU.

NOTE 10 When a system chooses to process these optional fields, the precise actions are not specified by this International Standard.

6.3.1 Record Configuration by Intermediate Systems

On receipt of an ESH PDU an IS extracts the configuration information and stores the {NSAP,SNPA} pairs in its local routeing information base replacing any other information for the same {NSAP,SNPA} pair. If insufficient space is available to store the new configuration information the PDU is discarded.

6.3.2 Record Configuration by End Systems

On receipt of an ISH PDU an ES extracts the configuration information and stores the {NET,SNPA} pairs in its local routeing information base replacing any other information for the same {NET,SNPA} pair. If insufficient space is available to store the new configuration information the PDU is discarded.

In addition, an ES may also recompute its Configuration Timer based on receipt of an ISH PDU containing the *Suggested ES Configuration Timer* (ESCT) optional field. If an End System chooses to use a computed CT rather than a local value supplied by System Management, it performs the operations described below.

- It examines its local routeing information base and ascertains whether any IS for which the ES is maintaining configuration information has supplied an ESCT. If no IS has suggested an ES configuration timer, the ES uses the value supplied by its local System Management.

- If one or more ISs suggested an ESCT, the minimum of the non-zero suggested values replaces the current value of the ES's CT.

6.4 Flush Old Configuration Function

The Flush Old Configuration function is executed to remove configuration entries in the routeing information base whose Holding Timer has expired. When the Holding Timer for an ES or IS expires, this function removes the corresponding entry from the routeing information base of the local Network entity.

The Flush Old Configuration function is also executed whenever a subnetwork service provider reinitializes a local SNPA. When the SNPA is either disabled or reinitialized, all configuration information for both ESs and ISs associated with that SNPA is removed.

6.5 Query Configuration Function

The Query Configuration function is performed under the following circumstances:

1) the End System is attached to a broadcast subnetwork,

2) there is no Intermediate System currently reachable on the subnetwork (i.e. no ISH PDUs have been received since the last information was removed by the Flush Old Configuration function),

3) the Network Layer's Route PDU function needs to obtain the SNPA address to which to forward a PDU destined for a certain NSAP,

4) the SNPA address cannot be obtained either by a local transformation or a local table lookup, and

5) QoS constraints would permit the broadcasting of the PDU.

NOTE 11 Despite appearances, this is actually a quite common case, since it is likely that there will be numerous isolated Local Area Networks without Intermediate Systems to rely upon for obtaining routeing information (e.g. via the Request Redirect Function of this International Standard). Further, if the Intermediate System(s) are temporarily unavailable, without this capability communication on the local subnetwork would suffer unless manually-entered tables were present in each End System or all NSAPs of the subnetwork had the subnetwork SNPA address embedded in them.

The End System, when needing to route an NPDU to a destination NSAP whose SNPA is unknown, issues an SN_UNITDATA.Request with the following parameters:

SN_Userdata ← NPDU

SN_Destination_Address ← multi-destination address that indicates "*All End System Network Entities*".

Subsequently an ESH PDU may be received containing the NSAP address along with the corresponding SNPA address (see 6.6). In such a case the End System executes the Record Configuration function for the NSAP, and therefore will be able to route subsequent PDUs to that destination using the specified SNPA. If no ESH PDU is received, the End System may declare the destination NSAP not reachable. The length of time to wait for a response before indicating a failure or the possibility of repeating the process some number of times before returning a failure are local matters and are not specified in this International Standard.

6.6 Configuration Response Function

The Configuration Response function is performed when an End System attached to a broadcast subnetwork receives an ISO 8473 PDU with the SN_Destination_Address from the SN_UNITDATA.Indication set to the multi-destination address "*All End System Network Entities*". This occurs as a result of another ES having performed the Query Configuration function described in 6.5.

If the ISO 8473 PDU is a valid PDU addressed to one of the NSAPS present in the ES, the End system shall:

1) Process the PDU according to the applicable clauses of ISO 8473;

2) Construct an ESH PDU containing information for at least that NSAP to which the received NPDU was addressed; and

3) Transmit the ESH PDU to the source of the original NPDU by issuing an SN_UNITDATA.Request with the following parameters:

   SN_Userdata ← ESH PDU

   SN_Destination_Address ← SN_Source_Address parameter value from the SN_UNITDATA.Indication containing the original NPDU as its SN_Userdata.

If the ISO 8473 PDU is either invalid or not addressed to one of the NSAPS present in the ES, the End system shall discard the PDU without generating an ISO 8473 Error Report.

6.7 Configuration Notification Function

The Configuration Notification function is used by End Systems and Intermediate Systems in order to transmit configuration information quickly to a system which has newly become available, in order to allow that system to build up its routeing information base as soon as possible.

A system which chooses to implement this function executes it on detecting, by receiving an ESH or ISH PDU, that another system has just become available. It then constructs an ISH or ESH PDU respectively, as described in 6.2.2 or 6.2.1, but transmits it specifically addressed to the newly operational system using an SN_UNITDATA.Request with the following parameters:

SN_Userdata ← ESH or ISH PDU

SN_Destination_Address ← SN_Source_Address parameter value from the SN_UNITDATA.Indication containing the original ESH or ISH PDU as its SN_Userdata.

It is recommended that systems which choose to implement this function should invoke it only when they can ascertain definitely that a system has recently become available and not, for example, simply because room for it has just become available in the routeing information base.

6.8 Request Redirect Function

The Request Redirect Function is present only in Intermediate Systems and is closely coupled with the Routeing and Relaying Functions of Intermediate Systems. The Request Redirect Function is coupled with the Route PDU Function described in ISO 8473. The Request Redirect Function is performed after the Route PDU function has calculated the next hop of the Data NPDU's path.

When an NPDU is to be forwarded by an Intermediate System, the Request Redirect Function first examines the out-

put of the IS's Routeing and Relaying function for this NPDU.

NOTE 12 As an optimization, the Request Redirect Function may examine the SN_Source_Address associated with the SN_UNITDATA.Indication which received the SNSDU (containing this NPDU). If it can be determined (for example by examining the configuration information obtained through the Record Configuration function) that the SN_Source_Address is not from an End System on the local subnetwork, then a Redirect PDU need not be sent.

This output will contain, among other things, the following pieces of information:

1) a local identifier for the subnetwork over which to forward the NPDU, plus either

2) the Network entity title and subnetwork address of the IS to which to forward the NPDU, or

3) the subnetwork address of the destination End System.

The Request Redirect function now determines if the source ES could have sent the NPDU directly to the Network entity the Intermediate System is about to forward the PDU to. Providing that QoS and other constraints permit NPDUs to by-pass this IS, then if any of the following conditions hold, the IS informs the source ES of the "better" path (by sending an RD PDU to the originating ES):

1) The next hop is to the destination system, and the destination is directly reachable (at subnetwork address BSNPA) on the source ESs subnetwork, or

2) The next hop is to an Intermediate System which is connected to the same subnetwork as the ES.

If the better path exists, the IS first completes normal processing of the received NPDU and forwards it. It then constructs a Redirect PDU (RD PDU) containing the Destination Address of the original NPDU, the subnetwork address of the better next hop (BSNPA), the Network entity title of the IS to which the ES is being redirected (unless the redirect is to the destination ES), a Holding Time (HT), QoS Maintenance, Priority, and Security options that were present in the Data NPDU (these are simply copied from the Data PDU). The HT is set to the value of the local Redirect Timer (RT). See Annex B for a discussion of how to choose the value of RT. If there are insufficient resources to both forward the original NPDU and to generate and send an RD PDU, the original NPDU shall be given preference.

The Request Redirect function may also be invoked by an IS when it receives a PDU addressed to an NSAP that is not reachable from this IS but to which the IS knows the first hop of a route from the source to the destination NSAP. In this case the IS shall first follow the procedures defined in ISO 8473 clauses 6.9 and 6.10 for discarding the PDU and generation of an error report. On completion of this procedure it shall inform the originating system of a route to the destination NSAP by sending an RD PDU.

Optionally, the IS may include information in the RD PDU indicating a larger population of NSAP addresses to which the same redirection information applies. There are two optional fields for this purpose: the *Address Mask* option and the *SNPA Mask* option. Their usage depends on the fact that NSAP addresses are represented using the preferred binary encoding, as specified in 7.3.2.

There are three permitted cases for including or excluding the masks. In the first case, both masks are absent. In this case, the RD PDU conveys information about one NSAP address only. The information reveals the system to which the IS is routeing the NPDU that provoked the RD PDU. That system could be another IS, or it could be the destination ES itself.

In the second case, the RD PDU contains an Address Mask but no SNPA Mask. In this case, the RD PDU conveys information about an equivalence class of NSAP addresses. The information reveals the system to which the IS sends NPDUs addressed to members of the class. If an ES receiving such an RD PDU decides to heed the mask, it may forward PDUs destined for members of the class to the system indicated in the RD PDU.

In the third case, the RD PDU contains both masks. As in the second case, the RD PDU conveys information about an equivalence class of NSAP addresses. But in this case, the information reveals that the SNPAs for that equivalence class of NSAPs are embedded in the NSAP. In particular, the SNPA Mask indicates the location of the SNPA in the NSAP. If an ES receiving such an RD PDU decides to heed the masks, it may route PDUs destined for members of the class directly to the SNPA extracted from the NSAP address.

The Intermediate System (assuming it has sufficient resources) then sends the RD PDU to the source End System by issuing an SN_UNITDATA.Reqeust with the following parameters:

SN_Userdata ← RD PDU

SN_Destination_Address ← SN_Source_Address parameter value from the SN_UNITDATA.Indication containing the original NPDU as its SN_Userdata.

6.9 Record Redirect Function

The Record Redirect Function is present only in End Systems. (ISs may receive RD PDUs, but do not process them). This function is invoked whenever an RD PDU is received. It extracts the redirect information and adds or replaces the corresponding redirection information in the local Network entity's routeing information base. The essential information is the redirection mapping from a Destination Address to a subnetwork address, along with the Priority, Security, and QoS Maintenance options and the Holding Time for which this mapping is to be considered valid. If the Redirect was to another Intermediate System, the Network entity title of the IS is recorded as well.

An End system may choose to ignore an RD PDU received for a destination to which the ES has not sent traffic for some period of time. An ES must record redirection information

only for tose other systems with which it is in active communication.

NOTE 13 If insufficient memory is available to store new redirection information, the RD PDU may be safely discarded since the original Intermediate System will continue to forward PDUs on behalf of this Network entity anyway.

6.10 Refresh Redirect Function

The Refresh Redirect Function is present only in End Systems. This function is invoked whenever an NPDU is received by a destination ES. It is closely coupled with the function that processes received NPDUs at a destination Network Entity (this is the PDU Decomposition function in ISO 8473). The purpose of this function is to increase the longevity of a redirection without allowing an incorrect route to persist indefinitely.

The Source Address (SA), Priority, Security, and QoS options are extracted and compared to any Destination Address and QoS parameters being maintained in the routeing information base (such information would have been stored by the Record Redirect Function). If a corresponding entry is found, the previous hop of the PDU is obtained from the SN_Source_Address parameter of the SN_Unitdata.Indication primitive by which it was received. If this address matches the next hop address stored with the redirection information, the remaining Holding Tmer for the redirection is reset to the original value that was obtained from the Holding Time field of the RD PDU. If the redirection information contains equivalence class masks, a separate Holding Timer is associated with this equivalence class information and is not reset.

NOTE 14 The purpose of this function is to avoid timing out redirection entries when the Network entity is receiving return traffic from the destination via the same path over which it is currently sending traffic. This is particularly useful when the destination system is on the same subnetwork as the source, since after one redirect no IS need be involved in the ES-to-ES traffic.

This function shall operate in a very conservative fashion however, to prevent the formation of black holes. The remaining holding timer shall be refreshed *only* under the exact conditions specified above. For a discussion of the issues surrounding the refresh of redirection information, see clause B.2 of Annex B.

6.11 Flush Old Redirect Function

The Flush Old Redirect Function is executed to remove redirection entries in the routeing information base whose Holding Timer has expired. When the Holding Timer expires, this function removes the corresponding entry from the routeing information base of the local Network entity.

The Flush Old Redirect function is also executed whenever a subnetwork service provider reinitializes a local SNPA. When the SNPA is either disabled or reinitialized, all redirection information associated with that SNPA is removed.

6.12 PDU Header Error Detection

The PDU Header Error Detection function protects against failure of Intermediate or End System Network entities due to the processing of erroneous information in the PDU header. The function is realized by a checksum computed on the entire PDU header. The checksum is verified at each point at which the PDU is processed. If the checksum calculation fails, the PDU is discarded.

The use of the Header Error Detection function is optional and is selected by the originating Network entity. If the function is not used, the checksum field of the PDU header is set to zero.

If the function is selected by the originating Network Entity, the value of the checksum field causes the following formulæ to be satisfied:

$$\sum_{i=1}^{L} a_i \pmod{255} = 0$$

$$\sum_{i=1}^{L} (L - i + 1) a_i \pmod{255} = 0$$

where L = the number of octets in the PDU header, and a_i = the value of the octet at position i. The first octet in the PDU header is considered to occupy position i=1.

When the function is in use, neither octet of the checksum field shall be set to zero.

6.13 Protocol Error Processing Function

A PDU in which the Network Layer Protocol Identifier field is present with the value defined in 7.2.2 and the Version/Protocol Identifier Extension is present with the value defined in 7.2.4, and which is not discarded by the PDU Header Error Detection function, shall be considered a protocol error if its encoding does not comply with the remainder of the provisions of 7. Any such protocol error PDU shall be discarded.

NOTE 15 PDUs in which the NPID has a value other than that in 7.2.2, or in which the V/P has a value other than in 7.2.4, are outside the scope of this International Standard.

7 Structure and Encoding of PDUs

NOTE 16 The encoding of the PDUs for this International Standard is compatible with that used in ISO 8473.

7.1 Structure

All Protocol Data Units contain an integral number of octets. The octets in a PDU are numbered starting from one (1) and increasing in the order in which they are put into an

SNSDU. The bits in an octet are numbered from one (1) to eight (8), where bit one (1) is the low-order bit.

When consecutive octets are used to represent a binary number, the lower octet number has the most significant value.

NOTE 17 When the encoding of a PDU is represented using a diagram in this section, the following representation is used:

1) octets are shown with the lowest numbered octet to the left, higher number octets being further to the right;

2) within an octet, bits are shown with bit eight (8) to the left and bit one (1) to the right.

PDU fields marked “Reserved” or “R” shall be treated as follows:

a) The value shall be set to binary zero by the transmitter, and

b) The value shall be ignored by the receiver.

PDUs contain, in the following order:

1) the Fixed part;
2) the Addressing Parameters part; and
3) the Options part, if present.

7.2 Fixed Part

7.2.1 General

The fixed part of the PDU header has the format shown in figure 1.

	Octet
Network Layer Protocol Identifier	1
Length Indicator	2
Version/Protocol Id Extension	3
Reserved	4
R R R Type	5
Holding Time	6,7
Checksum	8,9

Figure 1 - Fixed Part of PDU Header

7.2.2 Network Layer Protocol Identifier

The value of this field shall be 1000 0010.

This field identifies this Network Layer Protocol as ISO 9542.

7.2.3 Length Indicator

The length is indicated by a binary number, with a maximum value of 254 (1111 1110). The length indicated is the length of the entire PDU (which consists entirely of header, since this International Standard does not carry user data) in octets, as described in 7.1. The value 255 (1111 1111) is reserved for possible future extensions.

7.2.4 Version/Protocol Identifier Extension

The value of this field is binary 0000 0001. This identifies a standard version of ISO 9542.

7.2.5 Type Code

The Type code field identifies the type of the protocol data unit. The values defined for this field are given in table 2.

Table 2 - Valid PDU Types

	Bits	5	4	3	2	1
ESH PDU		0	0	0	1	0
ISH PDU		0	0	1	0	0
RD PDU		0	0	1	1	0

7.2.6 Holding Time

The Holding Time field specifies the maximum time for the receiving Network entity to retain the configuration/routeing information contained in this PDU.

The Holding Time field is encoded as an integral number of seconds.

7.2.7 PDU Checksum

The checksum is computed on the entire PDU header. A checksum value of zero is reserved to indicate that the checksum is to be ignored. The operation of the PDU Header Error Detection function (see 6.12) ensures that the value zero does not represent a valid checksum.

7.3 Addressing Parameters Part

7.3.1 General

Address parameters are distinguished by their location. The different PDU types carry different address parameters. The ESH PDU carries one or more Source NSAP addresses (SA); the ISH PDU carries an Intermediate System Network Entity Title (NET); and the RD PDU carries a Destination NSAP address (DA), a Subnetwork Address (BSNPA), and possibly a Network Entity Title (NET).

The address information is of variable length. Each address parameter is encoded as shown in figure 2.

Octet n	Address Parameter Length Indicator (e.g. 'k')
Octets $n+1$ thru $n+k$	Address Parameter Value

Figure 2 - Address Parameter Encoding

7.3.2 NPAI (Network Protocol Address Information) Encoding

The Destination and Source Addresses are Network Service Access Point addresses as defined in ISO 8348/Add.2. The Network Entity Title address parameter is a Network entity title as defined in ISO 8348/Add.2. The Destination Address, Source Address, and Network Entity Title are encoded as NPAI using the binary syntax defined in ISO 8348/Add.2.

7.3.3 Source Address Parameter for ESH PDU

The Source Address parameter is a list of one or more NSAP addresses of NSAPs served by the Network entity sending the ESH PDU. It is encoded in the ESH PDU as shown in figure 3.

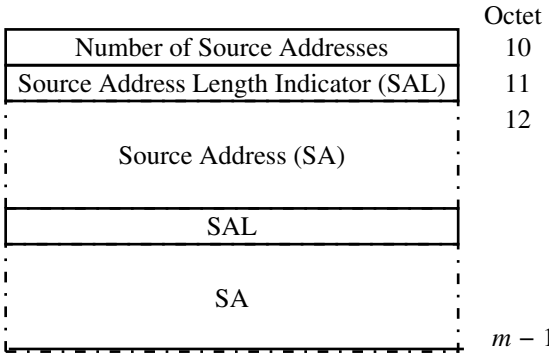

Figure 3 - ESH PDU — Source Address Parameter

7.3.4 Network Entity Title Parameter for ISH and RD PDUs

The Network Entity Title parameter is present in ISH and RD PDUs. In the ISH PDU, the NET parameter contains the Network entity title of the sending Intermediate system, as shown in figures 4 and 10. In the RD PDU, the NET parameter contains the Network entity title of the Intermediate system to which the End system is being redirected (figures 4 and 11), or contains an NETL field of zero as shown in figure 12.

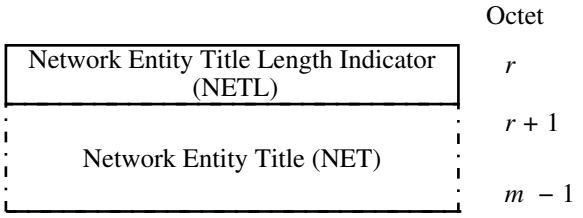

Figure 4 - ISH or RD PDU — Network Entity Title Parameter

7.3.5 Destination Address Parameter for RD PDU

The Destination Address is the NSAP address of a destination associated with some NPDU being forwarded by the Intermediate System sending the RD PDU. It is encoded in the RD PDU as shown in figure 5.

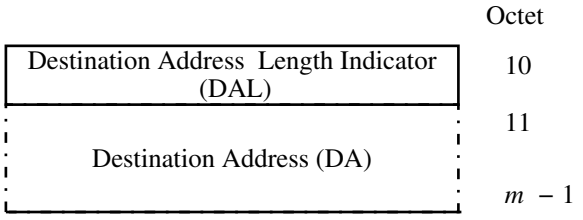

Figure 5 - RD PDU — Destination Address Parameter

7.3.6 Subnetwork Address Parameter for RD PDU

The Subnetwork Address Parameter is present only in RD PDUs. It is used to indicate the subnetwork address of another Network entity on the same subnetwork as the End System (and Intermediate System) which may be a better path to the destination specified in the Destination Address Parameter.

The Subnetwork Address Parameter is encoded in the RD PDU as shown in figure 6.

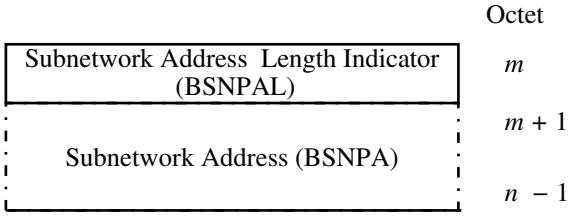

Figure 6 - RD PDU — Subnetwork Address Parameter

The Subnetwork Address is encoded in the form that is appropriate for the specific subnetwork to which it applies. In the case of an ISO 8802 subnetwork, the SNPA address is the

MAC address defined in ISO 10039, which is encoded according to the "binary representation of MAC addresses" specified in ISO 10039.

7.4 Options Part

7.4.1 General

The options part of the PDU header is illustrated in figure 8.

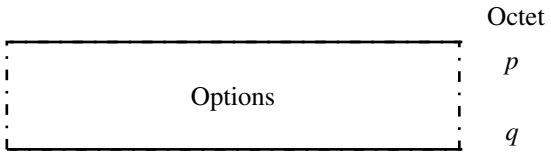

Figure 7 - All PDUs — Options Part

If the options part is present, it may contain one or more parameters. The number of parameters that may be contained in the options part is constrained by the length of the options part, which is determined by the following formula:

> PDU Header Length – (length of fixed part + length of addressing parameters part),

and by the length of the individual optional parameters.

If a PDU is received containing a parameter field whose parameter code is not listed in 7.4.2 onwards, that parameter field shall be ignored but the remainder of the PDU shall be processed as normal.

Parameters defined in the options part may appear in any order. Duplication of options is not permitted. Receipt of a PDU with an option duplicated shall be treated as a protocol error.

The encoding of parameters contained within the options part of the PDU header is illustrated below in figure 9.

Octets

n	Parameter Code
n + 1	Parameter Length
n+2 thru *n* + *m* + 1	Parameter Value

Figure 8 - Encoding of Option Parameters

The *parameter code field* is coded in binary and, without extensions, provides a maximum of 255 different parameters. No parameter codes use bits 8 and 7 with the value 00, so the actual maximum number of parameters is lower. A parameter code of 255 (binary 1111 1111) is reserved for possible future extensions.

The *parameter length field* indicates the length, in octets, of the parameter value field. The length is indicated by a positive binary number, *m*, with a theoretical maximum value of 254. The practical maximum value of *m* is lower. For example, in the case of a single parameter contained within the options part, two octets are required for the parameter code and the parameter length indicators. Thus, the value of *m* is limited to:

> *m* = 252 – (length of fixed part + length of addressing parameters part)

For each succeeding parameter the maximum value of *m* decreases.

The *parameter value field* contains the value of the parameter identified in the parameter code field.

7.4.2 Security

The Security option may appear in the ESH, ISH, or RD PDU.

When carried in an RD PDU, the Security parameter conveys information about the security requested in the Data PDU that caused the containing RD PDU to be generated. When carried in the ESH or ISH PDU, the Security parameter conveys security information about the transmitting system.

This parameter has the same encoding and semantics as the Security parameter in ISO 8473.

> **Parameter Code:** 1100 0101
> **Parameter Length:** variable
> **Parameter Value:** *See ISO 8473*

7.4.3 Quality of Service Maintenance

The QoS Maintenance option may appear only in the RD PDU.

The Quality of Service parameter conveys information about the quality of service requested in the Data PDU that caused the containing RD PDU to be generated.

This parameter has the same encoding and semantics as the QoS Maintenance parameter in ISO 8473.

> **Parameter Code:** 1100 0011
> **Parameter Length:** variable
> **Parameter Value:** *See ISO 8473*

7.4.4 Priority

The Priority option may appear in the ESH, ISH, or RD PDU.

When carried in an RD PDU, the Priority parameter conveys information about the priority requested in the Data PDU that caused the containing RD PDU to be generated. When carried in the ESH or ISH PDU, the Priority parameter conveys the priority of the transmitting system.

This parameter has the same encoding and semantics as the Priority parameter in ISO 8473.

> **Parameter Code:** 1100 1101

Parameter Length: one octet
Parameter Value: *See ISO 8473*

7.4.5 Address Mask

The Address Mask option may appear only in the RD PDU.

The Address Mask parameter indicates that the redirection information applies to a larger population of NSAP addresses than the Destination Address of the RD PDU indicates. An End System may ignore this parameter.

The Address Mask establishes an equivalence class of NSAP addresses to which the same redirection information applies. To determine whether or not a trial NSAP address falls within the equivalence class, the ES aligns the trial NSAP address with the Address Mask, padding the latter with trailing zero octets if necessary. If in all bit positions where the Address Mask is "1" the trial NSAP address matches the DA field of the RD PDU, then the trial NSAP address belongs to the equivalence class described by the RD PDU. In making routeing decisions, an exact NSAP address match takes precedence over use of equivalence classes. An exact match occurs when the trial NSAP address is identical to the one contained in the DA field of the RD PDU, without considering any mask. If a destination is within more than one equivalence class, the choice of which, if any, to use is a local matter.

An all zero Address Mask can be used to indicate an omniscient IS for outgoing NPDUs for which no route is otherwise known.

NOTE 18 By choosing an Address Mask according to the boundaries in the hierarchically administered NSAP address, the Address Mask permits routeing by subnetwork, by routeing domain, or by other administratively controlled criteria.

The Address Mask parameter has additional semantics when considered with the SNPA Mask parameter; see 7.4.6.

Parameter Code: 1110 0001
Parameter Length: variable, up to 20 octets
Parameter Value: a comparison mask of octets to be aligned with the Destination Address.

7.4.6 SNPA Mask

The SNPA Mask option may only appear in the RD PDU.

When the SNPA Mask is present, the equivalence class defined by the Address Mask also has common structure below the Address Mask; i.e. in the portion of the NSAP address where the Address Mask is logically "0". The SNPA Mask supplies additional information about that structure, by indicating certain bit positions within the space "below" the Address Mask. Specifically, the SNPA Mask indicates the location of the SNPA in the NSAP address.

This parameter may appear in an RD PDU only if the Address Mask is also present. An ES receiving such an RD PDU may safely ignore both masks. However (since presence of both masks dictates different functional behavior than the presence of the Address Mask alone) an ES shall not ignore one of the masks while heeding the other.

Parameter Code: 1110 0010
Parameter Length: variable
Parameter Value: a comparison mask of octets to be aligned with the Destination Address.

7.4.7 Suggested ES Configuration Timer

The ESCT option may appear only in the ISH PDU.

The ESCT parameter conveys the value that an IS would like the receiving ESs to use as their local Configuration Timer.

Parameter Code: 1100 0110
Parameter Length: two octets
Parameter Value: ESCT in units of seconds

7.5 End System Hello (ESH) PDU

The ESH PDU has the format shown in figure 10.

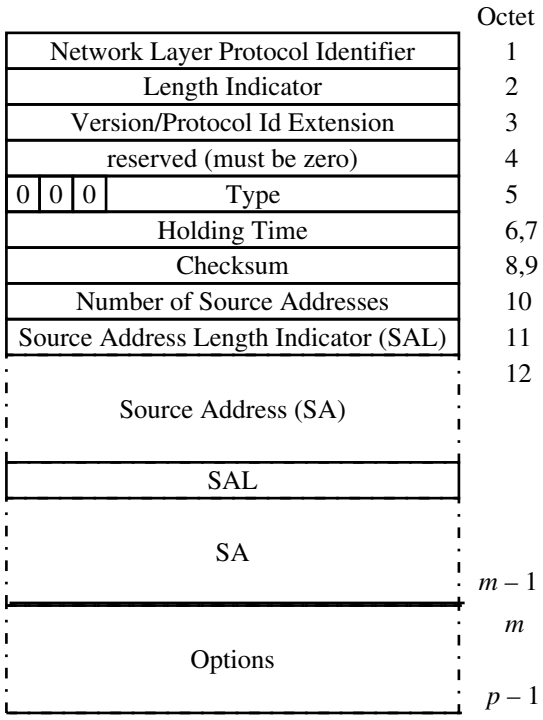

Figure 9 - ESH PDU Format

7.6 Intermediate System Hello (ISH) PDU

The ISH PDU has the format shown in figure 10.

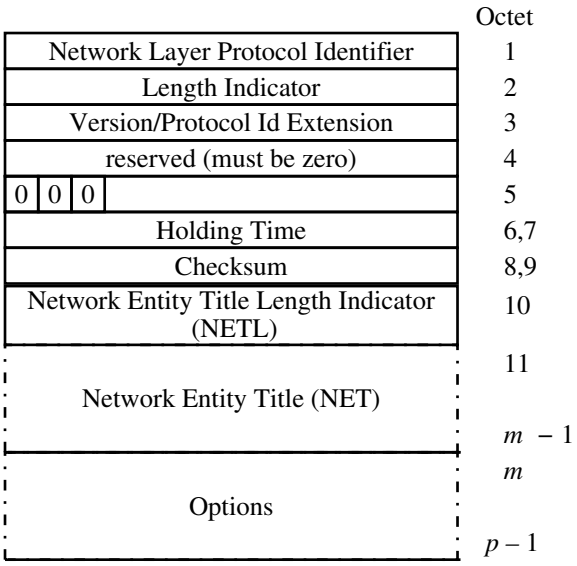

Figure 10 - ISH PDU Format

7.7 Redirect (RD) PDU

The RD PDU has the format shown in figures 12 and 13.

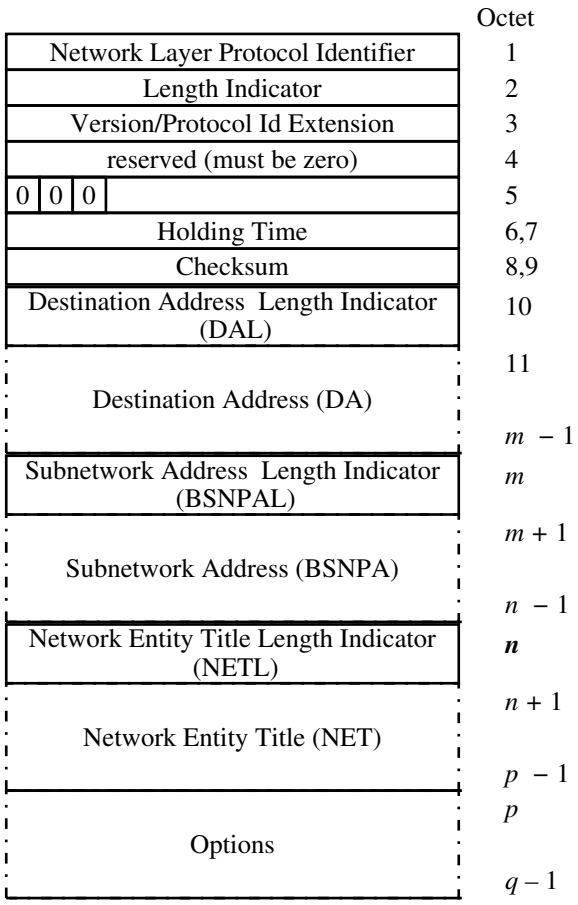

Figure 11 - RD PDU Format when redirect is to an IS

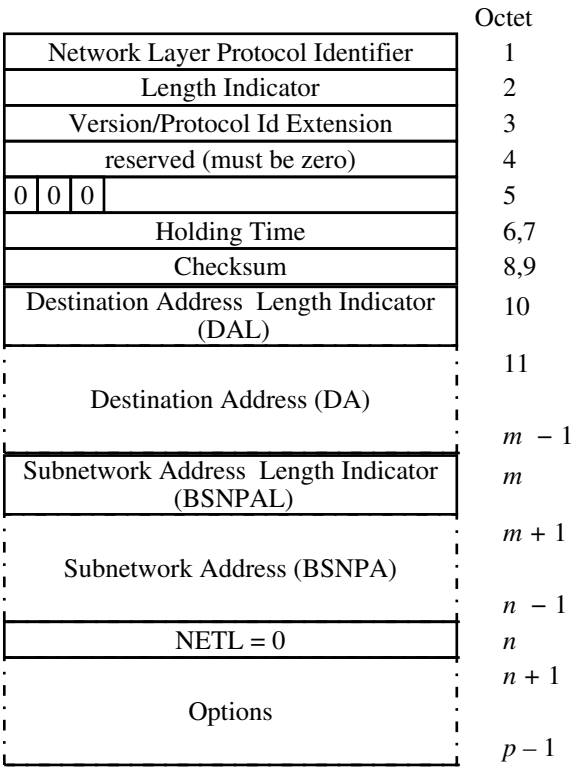

Figure 12 - RD PDU Format when redirect is to an ES

8 Conformance

8.1 Static Conformance Requirements

A Network entity may choose to support either the configuration information, the route redirection information, neither, or both. If the configuration information is supported, it is not required that it be employed over all subnetworks to which the Network entity is attached.

Implementations are not required to support all of the functions described in 6. Some functions are entirely optional, and the requirement for most of the remaining functions depends upon whether the implementation is for an End System or an Intermediate System, and upon whether the implementation supports configuration information, redirection information, both, or (for an ES only) neither. Table 3 and the following subclauses specify the requirements in the various cases.

8.2 Static Conformance Requirements for End Systems

An End System implementation that supports configuration information shall implement the functions marked as Mandatory (M) in column 5 of table 3.

An End System implementation that supports redirection information shall implement the functions marked as Mandatory (M) in column 6 of table 3.

An End System implementation that supports both configuration information and redirection information shall implement all the functions marked as Mandatory (M) in either column 5 or column 6 of table 3.

An End System implementation that supports neither configuration information nor redirection information shall implement the Configuration Response function, as marked Mandatory in column 4 of table 3.

8.2.1 Static Conformance Requirements for Intermediate Systems

An Intermediate System implementation that supports configuration information shall implement the functions marked as mandatory in column 7 of table 3.

An Intermediate System implementation that supports redirection information shall implement the functions marked as mandatory in column 8 of table 3.

An Intermediate System implementation that supports both configuration information and redirection information shall implement the functions marked as mandatory in either column 7 or column 8 of table 3.

8.3 Dynamic Conformance

Any protocol function supported shall be implemented in accordance with the appropriate subclause of 6.

Any PDU transmitted shall be constructed in accordance with the appropriate subclauses of 7.

8.4 Protocol Implementation Conformance Statement

A Protocol Implementation Conformance Statement (PICS) shall be completed in respect of any claim for conformance of an implementation to this International Standard: the PICS shall be produced in accordance with the relevant PICS proforma in Annex A.

Table 3 - Static Conformance Requirements

Label	Function	Defining Clause	ES			IS	
			min	CI	RI	CI	RI
ErrP	Protocol Error Processing PDU Header Error Detection	6.13	—	M	M	M	M
HCsV	• Checksum validation	6.12	—	M	M	M	M
HCsG	• Checksum generation	6.12	O	O	O	O	O
CfRs	Configuration Response	6.6	M	M	M	—	—
RpCf	Report Configuration	6.2	—	M	—	M	—
CTGn	• ESCT Generation	6.2.2	—	—	—	O	—
RcCf	Record Configuration	6.3	—	M	—	M	—
CTPr	• ESCT Processing	6.3.2	—	O	—	—	—
FlCf	Flush Old Configuration	6.4	—	M	—	M	—
QyCf	Query Configuration	6.5	—	M	—	—	—
CfNt	Configuration Notification	6.7	—	O	—	O	—
RqRd	Request Redirect	6.8	—	—	—	—	—
RcRd	Record Redirect	6.9	—	—	M	—	—
FlRd	Flush Old Redirect	6.11	—	—	M	—	—
RfRd	Refresh Redirect	6.10	—	—	O	—	—

Key:

CI = Configuration information supported
RI = Redirection information supported
min = neither supported (minimum ES implementation)
M = Mandatory O = Optional — = not applicable

Annex A
PICS Proformas

(Normative)

A.1 Introduction

The supplier of a protocol implementation which is claimed to conform to International Standard ISO 9542, whether as an End System or Intermediate System implementation, shall complete the applicable Protocol Implementation Conformance Statement (PICS) proforma following, and accompany it by the information necessary to identify fully both the supplier and the implementation.

A.2 Abbreviations and Special Symbols

A.2.1 General

N/A	Not Applicable
PR	receive (PDU, or field of PDU)
PS	send (PDU, or field of PDU)

A.2.2 Option-status and Predicate Symbols

M	mandatory
O	optional
P	prohibited
CI:	the status following this symbol applies only when the PICS states that configuration information is supported.
RI:	the status following this symbol applies only when the PICS states that redirection information is supported.
(CI∨RI):	the status following this symbol applies only when the PICS states that either configuration information or redirection information (or both) is supported.

A.3 Instructions for Completing the PICS Proformas

The main part of each PICS proforma is a fixed-format questionnaire. A supplier may also provide, or be required to provide, additional information, categorized as either Exception Information or Supplementary Information. When present, each kind of additional information is to be provided as items labelled respectively X.<i> or S.<i> for cross-referencing purposes, where <*i*> is any unambiguous identification for the item (e.g. simply a number): there are no other restrictions on its format and presentation.

A completed PICS proforma is the Protocol Implementation Conformance Statement for the implementation in question.

Answers to the questionnaires are to be provided in the rightmost column, either by simply marking an answer to indicate a restricted choice (such as *Yes* or *No*), or by entering a value or a set or range of values.

Items of Exception Information are required by certain answers in the questionnaire: this is indicated by an "X___" cross-reference to be completed. This occurs when, for example, an answer indicates that a feature classified as Mandatory has not been implemented: the Exception item should contain the appropriate rationale.

The final section of the PICS, for Supplementary Information, allows a supplier to provide additional information intended to assist the interpretation of the PICS. It is not intended or expected that a large quantity will be supplied, and a PICS can be considered complete without any such information. An example might be an outline of the ways in which a (single) implementation can be set up to operate in a variety of environments and configurations.

References to items of Supplementary Information may be entered next to any answer in the questionnaire, and may be included in items of Exception Information.

A.4 PICS Proformas

PICS Proforma: ISO 9542(1988) — End System

Item	Protocol Function	Clauses	Status	Support
CI	Is configuration information supported?		O	Yes No
RI	Is redirection information supported?		O	Yes No
	Are the following Functions supported?			
	Configuration Response			
CfRs	Protocol Error Processing	6.6	M	Yes No:X___
ErrP	PDU Header Checksum Validation	6.13	(CI∨PI):M	N/A Yes No:X___
HCsV	PDU Header Checksum Generation	6.12	(CI∨PI):M	N/A Yes No:X___
HCsG	Report Configuration	6.12	O	Yes No:X___
RpCf	Record Configuration	6.2, 6.2.1	CI:M	N/A Yes No:X___
RcCf	Flush Old Configuration	6.3, 6.3.2	CI:M	N/A Yes No:X___
FlCf	Query Configuration	6.4	CI:M	N/A Yes No:X___
QyCf	Record Redirect	6.5	CI:M	N/A Yes No:X___
RcRd	Flush Old Redirect	6.9	RI:M	N/A Yes No:X___
FlRd	Refresh Redirect	6.11	RI:M	N/A Yes No:X___
RfRd	Configuration Notification	6.10	RI:O	N/A Yes No
CfNt	ESCT Processing	6.7	CI:O	N/A Yes No
CTPr	Address Mask (only) Processing	6.3.2	CI:O	N/A Yes No
AMPr	Address Mask and SNPA Mask Processing	7.4.5	RI:O	N/A Yes No
SMPr		7.4.5, 7.4.6	RI:O	N/A Yes No

PICS Proforma: ISO 9542(1988) — End System (continued)

Item	Are the following PDUs supported?	Clauses	Status	Support
ESH-s	<s> End System Hello	7.1, 7.5	M	Yes No:X___
ESH-r	<r> End System Hello	7.1, 7.5	CI:M	N/A Yes No:X___
ISH-r	<r> Intermediate System Hello	7.1, 7.6	CI:M	N/A Yes No:X___
RD-r	<r> Redirect	7.1, 7.7	RI:M	N/A Yes No:X___
	Are the following PDU fields supported?			
	<s> Fixed Part			
FxPt	<r> Fixed Part	7.2.1–7.2.7	M	Yes No:X___
	<s> Source Address,	7.2.1–7.2.7	(CI∨RI):M	N/A Yes No:X___
SA-s1	<r> one NSAP only	7.3.1,	O.1	Yes No:X___
SA-r1	<s> Source Address,	7.3.2	CI:M	N/A Yes No:X___
SA-sm	<r> two or more NSAPs	7.3.3	O.1	Yes No:X___
SA-rm	<r> Network Entity Title		CI:M	N/A Yes No:X___
NET-r	<r> Destination Address	7.3.1, 7.3.2, 7.3.4	(CI∨RI):M	N/A Yes No:X___
DA-r	<r> Subnetwork Address	7.3.1, 7.3.2, 7.3.5	RI:M	N/A Yes No:X___
BSNPA-r	<s> Security	7.3.1, 7.3.2, 7.3.6	RI:M	N/A Yes No:X___
Scty-s	<r> Security	7.4.2	O	Yes No:X___
Scty-r	<s> Priority	7.4.2	O	Yes No:X___
Pty-s	<r> Priority	7.4.4	O	Yes No:X___
Pty-r	<r> QoS Maintenance	7.4.4	O	Yes No:X___
QoSM-r	<r> Address Mask	7.4.3	RI:O	N/A Yes No:
AdMk-r	<r> SNPA Mask	7.4.5	RI:O	N/A Yes No:
SNMk-r	<r> Suggested ES Configuration Timer	7.4.6	RI:O	N/A Yes No:
ESCT-r	<r> (ignore) unsupported or unknown options	7.4.7	CI:O	N/A Yes No:
OOpt-r	<s> Other options	7.4.1	M	Yes No:X___
OOpt-s			P	No Yes:X___
	Parameter Ranges			
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	From: seconds To: seconds by increments of[1]: (other — specify)[1]: with a tolerance of:
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	From: seconds To: seconds by increments of[1]: (other — specify)[1]: with a tolerance of:

[1] delete if inapplicable

PICS Proforma: ISO 9542(1988) — Intermediate System

Item	Protocol Function	Clauses	Status	Support
CI	Is configuration information supported?		O	Yes No
RI	Is redirection information supported?		O	Yes No
	Are the following Functions supported?			
	Protocol Error Processing			
ErrP	PDU Header Checksum Validation	6.13	M	Yes No:X___
HCsV	PDU Header Checksum Generation	6.12	M	Yes No:X___
HCsG	Report Configuration	6.12	O	Yes No
RpCf	Record Configuration	6.2, 6.2.2	CI:M	N/A Yes No:X___
RcCf	Flush Old Configuration	6.3, 6.3.1	CI:M	N/A Yes No:X___
FlCf	Request Redirect	6.4	CI:M	N/A Yes No:X___
RqRd	Configuration Notification	6.8	RI:M	N/A Yes No:X___
CfNt	ESCT Generation	6.7	CI:O	N/A Yes No
CTGn	Address Mask (only) Generation	6.3.2	CI:O	N/A Yes No
AMGn	Address Mask and SNPA Mask Genera-	6.8	RI:O	N/A Yes No
SMGn	tion	6.8	RI:O	N/A Yes No

PICS Proforma: ISO 9542(1988) — Intermediate System (continued)

Item	Are the following PDUs supported?	Clauses	Status	Support
ESH-r	<r> End System Hello	7.1, 7.5	CI:M	N/A Yes No:X___
ISH-r	<r> Intermediate System Hello	7.1, 7.6	CI:O	N/A Yes No
ISH-s	<s> Intermediate System Hello	7.1, 7.6	CI:M	N/A Yes No:X___
RD-s	<s> Redirect	7.1, 7.7	RI:M	N/A Yes No:X___
RD-r	<r> (ignore) Redirect	6.9, 7.1, 7.7	M	Yes No:X___
	Are the following PDU fields supported?			
FxPt	<s> Fixed Part	7.2.1–7.2.7	M	Yes No:X___
	<r> Fixed Part	7.2.1–7.2.7	M	Yes No:X___
SA-r	<s> Source Address, one or more NSAPs	7.3.1, 7.3.2, 7.3.3	CI:M	N/A Yes No:X___
NET-s	<s> Network Entity Title	7.3.1, 7.3.2, 7.3.4	M	N/A Yes No:X___
NET-r	<r> Network Entity Title	7.3.1, 7.3.2, 7.3.4	ISH-r::M	N/A Yes No:X___
DA-s	<s> Destination Address	7.3.1, 7.3.2, 7.3.5	RI:M	N/A Yes No:X___
BSNPA-s	<s> Subnetwork Address	7.3.1, 7.3.2, 7.3.6	RI:M	N/A Yes No:X___
Scty-s	<s> Security	7.4.2	O	Yes No
Scty-r	<r> Security	7.4.2	O	Yes No
Pty-s	<s> Priority	7.4.4	O	Yes No
Pty-r	<r> Priority	7.4.4	O	Yes No
QoSM-s	<s> QoS Maintenance	7.4.3	RI:O	N/A Yes No
AdMk-s	<s> Address Mask	7.4.5	RI:O	N/A Yes No
SNMk-s	<s> SNPA Mask	7.4.6	RI:O	N/A Yes No
ESCT-s	<s> Suggested ES Configuration Timer	7.4.7	CI:O	N/A Yes No
ESCT-r	<r> (ignore) Suggested ES Configuration Timer	7.4.7	ISR-r::M	N/A Yes No:X___
OOpt-r	<r> (ignore) unsupported or unknown options	7.4.1	M	Yes No:X___
OOpt-s	<s> Other options		P	No Yes:X___
	Parameter Ranges			
HTv	What range of values can be set for the Holding Time field in transmitted PDUs?	6.1, 6.1.2	M	From: seconds To: seconds by increments of[1]: (other — specify)[1]: with a tolerance of:
CTv	If configuration information is supported, what range of values can be set for the Configuration Timer?	6.1, 6.1.1	CI:M	From: seconds To: seconds by increments of[1]: (other — specify)[1]: with a tolerance of:

[1]delete if inapplicable

Annex B
Supporting Technical Material

(Informative)

B.1 Use of Timers

This International Standard makes extensive use of timers to ensure the timeliness and accuracy of information disseminated using the Configuration and Route Redirection functions. This section discusses the rationale for using these timers and provides some background for how they operate.

Systems using this International Standard learn about other systems exclusively by receiving PDUs sent by those systems. In a connectionless environment, a system must periodically receive updated information to ensure that the information it previously received is still correct. For example, if a system on a subnetwork becomes unavailable (either it has ceased operating, or its SNPA becomes inoperative) the only way another system can detect this fact is by the *absence* of transmissions from that system. If information were retained in the absence of new PDUs being received, configuration and/or routeing information would inevitably become incorrect. The Holding Timers specified by this International Standard guarantee that old information will not be retained indefinitely.

A useful way of thinking of the configuration and route redirection information is as a *cache* maintained by each system. The cache is periodically flushed to ensure that only up-to-date information is stored. Unlike most caches, however, the time to retain information is not a purely local matter. Rather, information is held for a period of time *specified by the source of the information*. Some examples will help clarify this operation.

B.1.1 Example of Holding Timer for Route Redirection

Route redirection information is obtained by an End System through the Request Redirect function (see 6.8). It is quite possible that an Intermediate System might redirect an End System to another IS which has recently become unavailable (this might happen if the IS-to-IS routeing algorithm is still converging following a configuration change). If the Holding Timer were not present, or was set very long by the sending IS, an End System would have been redirected into a {\large\bf Black Hole} from which none of its Data PDUs would ever emerge. The length of the Holding Time field in Redirect PDUs specifies, in essence, the length of time black holes are permitted to exist.

On the other hand, setting the Holding Time field on Route Redirects to a very small value to minimize the effect of black holes has other undesirable consequences. First, for each PDU that causes a redirect, an additional PDU besides the original Data PDU must be composed and transmitted; this increases overhead. Second, each time a "working" redirect's Holding Timer expires, the redirected End System will revert to a poorer route for at least one PDU.

B.1.2 Example of Holding Timer for Configuration Information

A similar type of problem can occur with respect to configuration information. If the Holding Time field of an ISH PDU (see \ref{isconfig}) is set to a very large value, and the only Intermediate System (which has been sending this configuration information) on the subnetwork becomes unavailable, a subnetwork-wide black hole can form. During this time, End Systems on the subnetwork may not be able to communicate with each other because they presume that an Intermediate System is operating which will forward their Data PDUs to destination ESs on the local subnetwork and return RD PDUs. Once the Holding Timer expires, the ESs will realize that no IS is available and will take their only recourse, which is to send their traffic directly on the local subnetwork.

Given the types of problems that can occur, it is important that responsibility for incorrect information can be unambiguously assigned to the source of the information. For this reason all Holding Timers are calculated by the source of the configuration or redirection information and communicated explicitly to each recipient in the Holding Time field of the appropriate PDU.

B.2 Refresh and timeout of redirection information

The protocol allows End Systems to refresh redirection information without first allowing the Holding Timer to expire and being redirected by a Intermediate System for a second (or subsequent) time. Such schemes are prevalent in connectionless subnetworks and are often called "reverse path information", "previous hop cache" or something similar.

Refreshing the redirection information has obvious performance benefits, but can be dangerous if not handled in a very conservative fashion. In order for a redirection to be safely refreshed, all of the following conditions must hold:

1) The source address of the received PDU must be exactly the same as the destination address specified in a prior RD PDU (this defines a "match" on the redirection information). Making assumptions about the equivalence of abbreviated addresses, multi-destination addresses, or similar "special" addresses is dangerous since routeing for these addresses cannot be assumed to be the same.

2) The Quality of Service parameters of the received PDU must be exactly the same as the QoS parameters specified in the matching (by destination address) redirection entry. Again, there is no guarantee that PDUs with different QoS parameters will be routed

the same way. It is quite possible that the redirected path is even a black hole for certain values of the QoS parameters (the security field is a good example).

3) The "previous hop" of the received Data PDU must match the "next hop" stored in the redirection information. Specifically, the SN_Source_Address of the SN_UNITDATA.Indication which received the PDU must match exactly the SN_Destination_Address specified in the redirect to be used for sending traffic via the SN_UNITDATA.Request primitive. This comparison ensures that redirects are refreshed only when the reverse traffic is being received from the same IS (or destination ES) as the forward traffic is being sent through (or to). This check make certain that redirects are not refreshed just on the basis of traffic being received from the destination. It is quite possible that the traffic is simply indicating that the forward path in use is not working!

Note that these conditions still allow refresh in the most useful and common cases where either the destination is another ES on the same subnetwork as the source ES, or the redirection is to an IS which is passing traffic to/from the destination in both directions (i.e. the path is symmetric).

B.3 System Initialization Considerations

This International Standard is designed to make the exchange of information as free as possible from dependencies between the two types of systems. Therefore, it is not possible for an End System to request all Intermediate Systems on a subnetwork to report their configuration, nor is it possible for an Intermediate System to request all End Systems on a subnetwork to report their configuration.

In certain operating environments a constraint may be imposed than an ES, upon becoming operational, must discover the existence of an IS as soon as possible. The converse relationship also holds if it is necessary for an IS to discover the existence of End Systems as soon as possible. In both cases the availability of this information is normally determined by the Configuration Timer of the system for which the knowledge is desired. There is therefore a tradeoff between the overhead associated with performing the Report and Record Configuration functions and the timely availability of the configuration information. Decreasing the Configuration Timer increases the availability at the expense of an increase in overhead.

The following solution is accommodated in the standard in the following way. When the Record Configuration function is invoked in either an End System or an Intermediate System, the function will determine if the received configuration information was previously unknown. If this is the case, then the Report Configuration function may be invoked before the expiration of the system's Configuration Timer. The Hello PDU generated by the Report Configuration function is then sent *only* to the Network Entity whose configuration was previously unknown. Thus when an ES or IS first becomes operational it immediately reports its configuration. As soon as systems of the other type discover the new network entity, they will make their own configurations known to this entity.

The additional overhead incurred by this solution is minimal. Also, since the discovery of new configurations is made timely by this approach the Configuration Timer period can be increased in order to decrease the overhead of the configuration functions, provided that other factors not discussed here are accounted for by the longer time period. One caveat is that the first Hello PDU generated by a system may be lost during transmission. To solve this problem one or more additional PDUs may be transmitted at short time intervals during this initialization period.

B.4 Optimizations for Flushing Redirects

An ES will attempt to forward NPDUs through an IS to which it has been redirected until the Holding Timer specified in the Holding Time field of the RD PDU has expired, even if that IS is no longer reachable. Under certain circumstances, it is possible to do better and recognize the existence of a black hole sooner. In particular, if the ES expects to hear ISH PDUs from the IS to which it has been redirected, and the Holding Timer for that IS expires, all knowledge of the IS may be forgotten by the ES. This includes any redirects, which may be flushed (see the Flush Old Redirect function) even though their timeouts have not expired.

B.5 An Example of using the Address and SNPA Masks

Consider an NSAP address authority which decomposes the DSP into *k* hierarchical elements as shown in figure 13. When an ES receives an RD PDU containing an Address Mask, the mask identifies a matching pattern, or common structure, against which any destination NSAP addresses subsequently made known to the ES can be compared. The comparison determines whether the NSAP address is reachable by the information conveyed in the RD PDU.

With only one mask present, that information is the identity of the IS to use as the first immediate destination.

When, in addition, the SNPA Mask is contained in the RD PDU. This second mask identifies that hierarchical element within the NSAP address upon which the ES can apply its local routeing algorithm.

An example of how this general model might be applied is a private network whose topology consists of LANs, ISO8208-based wide area networks, and point-to-point HDLC-based links. In figure 14 ISs *A*, *B*, *C*, and *D* unite four logical groupings of ESs (1 through 4). The DSP used for such a topology might have two hierarchical levels, which (for lack of better names) are called {subnet-id,element}. The two masks for such an NSAP address might be defined as in figure 15

Applying the model, an ES on subnet 3 receives an RD PDU from IS *C* containing an Address Mask indicating that all destinations {4,*i*} are reachable via IS *D*. The ES records the fact that the specific NSAP can be reached via IS *D*, and also records the Addres Mask. When a request is made to send an NPDU to another (different) destination NSAP, the ES will compare the NSAP address with the Address Mask. If the corresponding bits of the destination NSAP address are equal to those of the recorded NSAP address, the ES forwards the NPDU to IS *D*. If the comparison fails (in this limited scenario), whether the ES forwards the NPDU to IS *C* or IS *D* is a local choice; if the choice is sub-optimal, it will be corrected by a subsequent RD PDU.

If an SNPA Mask accompanies the Address Mask in the RD PDU, the ES first applies the Address Mask comparison as before. This time, if the comparison yields a match, the ES uses the SNPA Mask to extract from the NSAP address the routeing information it uses in its local routeing function. In this scenario, the process would result in the isolation of the logical "element" number which is used as input to the ES's Route PDU function.

The particular leaf-level routeing method depends on the particular situation. One example is an address administration authority that chooses to embed the SNPA in the NSAP address for most of the ESs on a LAN subnetwork. In this example, the equivalence class indicates the set of ESs on that subnetwork for whom this convention is used, and the SNPA Mask indicates the position of the SNPA in the NSAP address. Other ESs on the same subnetwork can be administered individually without embedding the SNPA in the NSAP address. Still others on the same subnetwork can comprise a separate equivalence class set up by the rules of a different address administration authority. A second example deals with ESs accessible through that network, and the SNPA Mask indicates the location in the NSAP address of the X.121 address of the first (or only)hop in the path. A final example deals with an ES connected to a set of dedicated point-to-point links, as well as to other subnetworks. Some of the point-to-point links reach ESs; some reach ISs. The equivalence class indicates the set of ESs to be reached over the point-to-point links, and the SNPA Mask indicates where in the NSAP address to find an identifier from which the identity of one particular link can be derived.

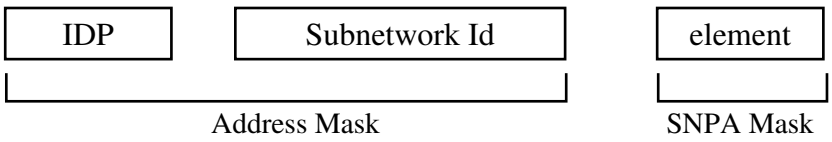

Figure 15 - Example Address and SNPA Masks

Annex C
State Tables

(Informative)

This Annex contains a summary of the protocol. It is provided to assist implementers of the protocol. In the event of a discrepancy between the description in these tables and that contained in the text, the text takes precedence.

This Annex describes the protocol in terms of state tables, tables 7 and 8, which show the state of an End System and an Intermediate System, the events that occur in the protocol, the actions taken, and the resultant state. Tables 4, 5 and 7, and the following descriptive text, define the notation used in the state tables themselves.

C.1 States

The only state defined for this International Standard is the **READY** state. While in the **READY** state, the protocol is capable of performing any of the functions enumerated in 6.

C.2 Events

Events are represented by their abbreviated names, as defined in table 4.

Table 4 - Events

Name	Description
CT	Configuration Timer expiry
HT_i	Holding Timer expiry for Table Entry$_i$
ESH	End System Hello PDU received
ISH	Intermediate System Hello PDU received
DT	ISO 8473 Data PDU received
RD	Redirect PDU received
ERR	Erroneous PDU received: 6.12, 6.13
SNPA	Local SNPA disabled or reinitialized
QC	Query Configuration needed: 6.5
ER	ISO 8473 Error PDU received

C.3 Actions and Predicates

For each event, the tables specify a resultant set of actions: the actions are represented by abbreviated names as defined in table 7.

Some action-sets are conditional: this is indicated by an entry in the predicate column of the table. Predicates are boolean expressions using abbreviated names defined in table 5. An action-set is only performed if its corresponding predicate is true. The predicate **false** in table 8 indicates that the QC event never applies to ISs.

C.4 Addressing Notation

In table 5, *SN-DA* means the SN_Destination_Address parameter value of the SN_Unitdata.Indication primitive that contained the received Data PDU as its SN_Userdata parameter. Similarly, in table 7, *SN-SA* means the SN_Source_Address parameter of the relevant SN_Unitdata.Indication primitive. Also in table 7 *SN-DA* means the SN_Destination_Address parameter value of the SN_Unitdata.Request primitive generated to convey the transmitted PDU as its SN_Userdata.

The multi-destination addresses for “All End System Network entities” and “All Intermediate System Network entities” are abbreviated as *all-ESs* and *all-ISs* respectively.

Table 5 - Predicates

Name	Description
Pc	System supports configuration information
Pr	System supports redirection information
Pp	System elects to do prompt reporting of configuration on SNPA reinitialization: 6.2
Pe	End System elects to process ESH PDUs
Pi	Intermediate System elects to process ISH PDUs
P1	Insufficient space available to record configuration information
P2	Data PDU received with SN-DA = **all-ESs**
P3	Matching redirection information exists for originator of DT PDU, and system supports Refresh Redirect function
P4	Received ESH or ISH PDU is from a newly available system, and Configuration Notification function is supported
P5	A better path exists

Table 6 - Specific Actions

Name	Description
CT-reset	Stop and restart Configuration Timer
HT_i-reset	Stop and restart Holding Timer for table entry$_i$
$Flush_i$	Flush all configuration and/or redirection information for table entry$_i$
Record	Record new configuration and/or redirection information
Px-ESCT	Optionally, process ESCT parameter if present in PDU
Discard	Discard PDU
ESH:IS*	Send End System Hello PDU with SN-DA = **all ISs**
ESH:SN-SA	Send End System Hello PDU with SN-DA = received SN-SA
ISH:ES*	Send Intermediate System Hello PDU with SN-DA = **all ESs**
ISH:SN-SA	Send Intermediate System Hello PDU with SN-DA = received SN-SA
RD:SN-SA	Send Redirect PDU with SN-DA = received SN-SA
DT:ES*	Send Data PDU with SN-DA = **all ESs**
Forward-DT	Forward Data PDU as specified by ISO 8473
Forward-ER	Forward Error PDU as specified by ISO 8473

Table 7 - End System State Table

Event	Predicate	Actions	Clause	New State
CT	*Pc*	ESH:IS* ; CT-reset	6.2.1	READY
HT_i	*Pc* ∨ *Pr*	$Flush_i$	6.4, 6.11	
ISH	*Pc* ∧ *P1*	Discard	6.3, 6.3.2	
	Pc ∧¬ *P1* ∧¬*P4*	Record ; HT_i-reset ; Px-ESCT	6.3, 6.3.2	
	Pc ∧¬ *P1* ∧ *P4*	Record ; HT_i-reset ; Px-ESCT ; ESH:SN-SA	6.3, 6.3.2, 6.7	
ESH	*Pc* ∧ *Pe* ∧ *P1*	Discard	6.3, 6.3.2	
	Pc ∧¬ *P1* ∧ *Pe* ∧¬ *P4*	Record ; HT_i-reset	6.3, 6.3.2	
	Pc ∧¬ *P1* ∧ *P*ε ∧ *P4*	Record ; HT_i-reset ; ESH:SN-SA	6.3, 6.3.2, 6.7	
DT	*P2*	ESH:SN-SA	6.6	
	¬ *P2* ∧ *P3* ∧ *Pr*	HT_i-reset	6.10	
RD	*Pr* ∧ *P1*	Discard	6.9	
	Pr ∧¬ *P1*	Record; HT_i-reset	6.9	
ERR		Discard	6.12, 6.13	
QC	*Pc*	DT:ES*	6.5	
SNPA	*Pr* ∨ (*Pc* ∧¬ *Pp)*	Flush-all	6.4, 6.11	
	Pc ∧ *Pp*	Flush-all ; ESH:IS* ; CT-reset	6.4, 6.11, 6.2, 6.2.1	

Table 8 - Intermediate System State Table

Event	Predicate	Actions	Clause	New State
CT	*Pc*	ISH:IS* ; CT-reset	6.2.2	**READY**
HT_i		$Flush_i$	6.4	
ESH	*Pc* ∧ *P1*	Discard	6.3, 6.3.1	
	Pc ∧¬ *P1* ∧¬ *P4*	Record ; HT_i-reset	6.3, 6.3.1	
	Pc ∧¬ *P1* ∧ *P4*	Record ; HT_i-reset ; ESH:SN-SA	6.3, 6.3.1, 6.7	
ISH	*Pc* ∧ *Pi* ∧ *P1*	Discard	6.3, 6.3.1	
	Pc ∧¬ *P1* ∧ *Pi* ∧¬*P4*	Record ; HT_i-reset	6.3, 6.3.1	
	Pc ∧¬ *P1* ∧ *Pi* ∧ *P4*	Record ; HT_i-reset ; ISH:SN-SA	6.3, 6.3.1, 6.7	
DT	*Pr* ∧ *P5*	Forward-DT ; RD:SN-SA	6.8	
	¬ *P5*	Forward-DT	6.8	
ER	*Pr* ∧ *P5*	Forward-ER ; RD:SN-SA	6.8	
	¬ *P5*	Forward-ER	6.8	
RD		Discard	6.9	
ERR		Discard	6.12, 6.13	
QC	**false**	—		
SNPA	¬ *Pp*	Flush-all	6.4	
	Pc ∧ *Pp*	Flush-all ;ISH:ES* ; CT-reset	6.4, 6.2, 6.2.2	