

Open Systems Interconnection: Specification for protocol to provide the connectionless-mode network service

UDC 681.324.02.071.8

NO COPYING WITHOUT BSI PERMISSION EXCEPT AS PERMITTED BY COPYRIGHT LAW



Document Sharing Hub

Committees responsible for this British Standard

The preparation of this British Standard was entrusted by the Information Systems Technology Standards Policy Committee (IST/-) to Technical Committee IST/6, upon which the following bodies were represented:

- Association for Payment Clearing Services
- British Computer Society
- British Telecommunications plc
- Business Equipment and Information Technology Association
- Department of Trade and Industry (Information Technology Standards Unit)
- Department of Trade and Industry (National Physical Laboratory)
- Electricity Supply Industry in England and Wales
- Electronic Engineering Association
- HM Treasury (Central Computer and Telecommunications Agency)
- Institute of Data Processing Management
- Inter-Universities Computing Committee
- Joint Network Team
- National Computing Centre Ltd.
- National Health Service
- OFTEL (Office of Telecommunications)
- Post Office
- Telecommunication Engineering and Manufacturing Association
- Telecommunications Managers' Association

This British Standard, having been prepared under the direction of the Information Systems Technology Standards Policy Committee, was published under the authority of the Board of BSI and comes into effect on 30 April 1990

© BSI 04-2000

The following BSI references relate to the work on this standard:
Committee reference IST/6
Draft for comment DD 118

ISBN 0 580 18100 6

Amendments issued since publication

Amd. No.	Date of issue	Comments
7207	September 1992	Indicated by a sideline in the margin

Contents

	Page
Committees responsible	Inside front cover
National foreword	ii
0 Introduction	1
1 Scope and field of application	1
2 References	2
Section 1. General	
3 Definitions	3
4 Abbreviations	4
5 Overview of the protocol	6
Section 2. Specification of the protocol	
6 Protocol functions	9
7 Structure and Encoding of PDUs	17
8 Provision of the Underlying Service	29
9 Conformance	35
Annex A Formal Description	37
Annex B Supporting Technical Material	51
Annex C Algorithms for PDU Header Error Detection function	55
Figure 1 — Interrelationship of standards	1
Figure 2 — PDU Structure	17
Figure 3 — PDU Header — Fixed Part	18
Figure 4 — PDU Header — Address Part	20
Figure 5 — Address Parameters	20
Figure 6 — PDU Header — Segmentation Part	21
Figure 7 — PDU Header — Options Part	21
Figure 8 — Encoding of Option Parameters	21
Figure 9 — PDU Header — Data Part	25
Figure 10 — DT PDU	26
Figure 11 — Inactive Network Layer Protocol	26
Figure 12 — Error Report PDU	27
Table 1 — Connectionless-mode Network-Service primitive	7
Table 2 — Underlying Service primitive	7
Table 3 — Timer primitive	8
Table 4 — Categorization of Protocol functions	18
Table 5 — PDU Type Codes	19
Table 6 — Error Report PDU: Fixed Part Fields	27
Table 7 — Reason for Discard Parameter Values	28
Table 8 — ISO 8802-2 LLC Service primitives	32
Table 9 — Subnetwork Connection Reference Encoding	34
Table 10 — Provision of functions for Conformance	36
Publications referred to	Inside back cover

National foreword

This British Standard has been prepared under the direction of the Information Systems Technology Standards Policy Committee. It is identical with ISO 8473:1988, “*Information processing systems — Data communications — Protocol for providing the connectionless-mode network service, including Addendum 3*”, published by the International Organization for Standardization (ISO). It supersedes DD 118 which is withdrawn.

The UK has actively participated, through BSI Technical Committee IST/6, in the development of the International Standard.

Up-to-date information on existing and projected British Standards publications for Open Systems Interconnection is included in the “Basic List”, published as part of the BSI Information Technology Service (BITS) initiative. For details contact: Marketing Department, BSI, Linford Wood, Milton Keynes, MK14 6LE.

Cross-references

International Standard	Corresponding British Standard
ISO 7498:1984	BS 6568 <i>Reference model of open systems interconnection</i>
ISO 7498/Add. 1:1987	Part 1:1988 <i>Basic reference model (incorporating connectionless-mode transmission)</i> (Identical)
ISO 8208:1987	BS 7219:1989 <i>Specification for X.25 packet level protocol for data terminal equipment for information processing systems</i> (Identical)
ISO 8348:1987	^a BS 7220:1989 <i>Open Systems Interconnection: Network service definition (incorporating connectionless-mode transmission and network layer addressing)</i>
ISO 8348/Add. 1:1987	(Identical)
ISO 8348/Add. 2:1988	(Identical)
ISO/TR 8509:1987	DD 108:1988 <i>Open Systems Interconnection Service conventions (ISO/TR 8509)</i> (Identical)
ISO 8648:1988	BS 7221:1989 <i>Open Systems Interconnection: Specification for internal organization of the Network Layer</i> (Identical)

^a In preparation.

The Technical Committee has reviewed the provisions of ISO 8802, ISO 8886 and ISO 9074, to which reference is made in the text, and has decided that they are acceptable for use in connection with this standard.

A British Standard does not purport to include all the necessary provisions of a contract. Users of British Standards are responsible for their correct application.

Compliance with a British Standard does not of itself confer immunity from legal obligations.

Summary of pages

This document comprises a front cover, an inside front cover, pages i and ii, the ISO/IEC title page, pages ii to iv, pages 1 to 56, an inside back cover and a back cover.

This standard has been updated (see copyright date) and may have had amendments incorporated. This will be indicated in the amendment table on the inside front cover.

INTERNATIONAL STANDARD

ISO/IEC
8473/
AD3

First edition 1989-02-28



INTERNATIONAL ORGANIZATION FOR STANDARDIZATION
ORGANISATION INTERNATIONALE DE NORMALISATION

Information processing systems - Telecommunications and Information exchange between systems - Protocol for providing the connectionless-mode network service

ADDENDUM 3: Provision of the underlying service assumed by ISO 8473 over subnetworks which provide the OSI data link service

Systèmes de traitement de l'information - Communications de données - Protocole fournissant le service de réseau en mode sans connexion

ADDTIF 3: Fourniture du service sous-jacent assuré par l'ISO 8473 sur des sous-réseaux point à point fournissant le service de liaison de données OSI

Reference number

ISO/IEC 8473/AD3: 1989 (1)

Introduction

The Protocol for Providing the Connectionless-mode Network Service (ISO 8473) is one of a set of International Standards produced to facilitate the Interconnection of open systems. This Addendum 3 to ISO 8473 defines the provision of the underlying service assumed by the protocol described in ISO 8473 over subnetworks which provide the OSI Data Link Service.

The material in this Addendum will be incorporated into the body of ISO 8473 when the next revision of the Standard is produced. This Addendum has a structure which is similar to that of ISO 8473 and previous Addenda in order to facilitate cross reference between the documents and their eventual incorporation into ISO 8473.

1 Scope

This addendum to ISO 8473 define a set of Subnetwork Dependent Convergence Function (SNDCFs) that relate the OSI Data Link Service defined in ISO 8886 to the abstract “underlying service” defined in subclause 5.5 of ISO 8473. The SNDCFs defined in this Addendum complement the set of SNDCFs already defined in clause 8, Provision of the Underlying Service.

The scope of this addendum is strictly limited to configurations in which the ISO 8473 protocol is used in conjunction with one or more subnetworks to provide the OSI Connectionless-mode Network Service. When the ISO 8473 protocol is used for this purpose, the appropriate subnetwork-dependent convergence functions defined in this addendum are applied to each subnetwork that participates in the transmission of ISO 8473 protocol data units from a source Network Service Access Point to a destination Network Service Access Point.

The field of application of this addendum is the operation of the ISO 8473 protocol over subnetworks which provide the OSI Data Link Service. For certain types of subnetworks, a circuit must be opened between end-points before communication can begin. The procedures for opening and closing circuits are defined in the standards pertaining to the particular subnetwork. The choice of when to open a circuit and to which end-point the circuit is opened are a local matter. This addendum defines the convergence functions to be used once the circuit exists.

This addendum affects clause 8 only and makes no additions or changes to the following:

- a) clause 2;
- b) clauses 3, 4 and 5 of Section 1;
- c) clauses 6, 7 and 9 of Section 2;
- d) annexes.

8 Provision of the Underlying Service

This clause complements clause 8 of ISO 8473.

8.1 Subnetwork Points of Attachment

This addendum makes no additions or changes to clause 8.1 of ISO 8473.

8.2 Subnetwork Quality of Service

Subclause 8.2.1 note 4): replace the text of note 4) with the following:

“In the case of an underlying subnetwork which provides a connection-mode service, the Transit Delay value provided to the CLNP must take into consideration any processing or queueing delay incurred as a result of an attempt to establish a connection.”

8.3 Subnetwork User Data

This addendum makes no additions or changes to clause 8.3 of ISO 8473.

8.4 Subnetwork Dependent Convergence Functions

This clause complements clause 8.4 of ISO 8473.

8.4.1 General Model

This addendum makes no additions or changes to clause 8.4.1 of ISO 8473.

8.4.2 Subnetwork Dependent Convergence Functions used with ISO 8802-2 Subnetworks

This addendum makes no additions or changes to clause 8.4.2 of ISO 8473.

8.4.3 Subnetwork Dependent Convergence Functions used with ISO 8208 Subnetworks

This addendum makes no additions or changes to clause 8.4.3 of ISO 8473.

8.4.4 Subnetwork Dependent Convergence Functions used with ISO 8886

Clause 8.4.4 is to follow clause 8.4.3 of ISO 8473.

This clause defines a mapping of the OSI Data Link Service to the underlying service assumed by ISO 8473. The OSI Data Link Service Definition defines two types of Data Link Services: a connectionless-mode service and a connection-mode service. SNDCFs are defined for subnetworks which provide either of these two modes of service.

8.4.4.1 *SNDCF used with the Connectionless-mode Data Link Service*

The primitives defined for provision of the underlying service assumed by the CLNP map directly onto the UNITDATA Request and Indication primitives defined for the Connectionless-mode Data Link Service. Subnetwork Dependent Convergence Functions perform a mapping of the Connectionless-mode Data Link Service onto the underlying service assumed by the CLNP. The mapping is as follows. The generation of an SN-UNITDATA request by the CLNP results in the generation of a DL-UNITDATA request (as described in ISO 8886) by the Subnetwork Dependent Convergence Function. A corresponding DL-UNITDATA indication prompts the SNDCF to generate an SN-UNITDATA indication to the CLNP. No explicit Subnetwork Dependent Convergence Protocol Control Information is exchanged between Network entities to provide this mapping of service.

The parameters of the SN-UNITDATA primitives are mapped onto the DL-UNITDATA primitives as follows. The SN-Destination-Address and SN-Source-Address parameters are conveyed in the DL-Destination-Address and DL-Source-Address parameters, respectively. The addresses used in the SN-UNITDATA request and indication primitives are the Data Link Service Access Point addresses described in ISO 8886.

The SN-Quality-of-Service parameter is conveyed. The available QoS is known prior to the issuance of the DL-UNITDATA.Request. There is no discrimination among DLSUs.

The SN-Userdata parameter is conveyed in the DL-Userdata parameter. The subnetwork must be able to support the service data unit requirements defined in clause 8.3.

8.4.4.2 *SNDCF used with the Connection-mode Data Link Service*

The primitives defined for provision of the underlying service assumed by the CLNP are mapped onto the primitives defined for the Connection-mode Data Link Service. Subnetwork Dependent Convergence Functions perform a mapping of the Connection-mode Data Link Service onto the underlying service assumed by the CLNP. The mapping is as follows.

On receiving an SN-UNITDATA request from the CLNP machine, the SNDCF determines if a Data Link Connection already exists between this source and destination pair to convey the user data. If so, it issues a DL-DATA request containing the SN-Userdata as the DL-Userdata.

If a Data Link Connection does not already exist, a DL-CONNECT request is issued by the local/calling SNDCF with the Source and Destination Data Link Service Access Point Addresses specified in the SN-UNITDATA request and waits for a DL-CONNECT confirm. The SN-Quality-of-Service parameter is conveyed. The available QoS is known prior to the issuance of the DL-CONNECT.Request. There is no requirement to use Expedited Data.

When the remote/called SNDCF receives a DL-CONNECT indication from the Data Link layer, it issues a DL-CONNECT response. Once the corresponding DL-CONNECT confirm is received by the local/calling SNDCF, it may then issue a DL-DATA request(s) conveying user data. When the remote/called SNDCF receives a DL-DATA indication from the Data Link layer, it issues an SN-UNITDATA indication which conveys the corresponding destination and source addresses as well as the SN-Userdata. The SN-Userdata parameter is conveyed in the DL-Userdata parameter. The subnetwork must be able to support the service data unit requirements defined in clause 8.3.

The mechanism and timing for opening a Data Link Connection prior to the transmission of SN-Userdata are a local matter. The opening of a Data Link Connection may be initiated by:

- a) the arrival of an SNSDU to be transmitted over a subnetwork at a time when no Data Link Connection is available;
- b) the local queue of requests waiting for an existing Data Link Connection reaching a threshold size at which an additional Data Link Connection must be made available (if possible) to maintain the requested QoS; or
- c) the explicit intervention of the network management system.

Collision detection and correction are resolved within the Data Link layer.

When it has been determined that a Data Link Connection must be cleared, the local/calling SNDCF issues a DL-DISCONNECT request primitive, specifying itself (the local DLS User) as the originator of the release, and a reason code as defined in ISO 8886. Once the request primitive has been issued, the local SNDCF considers the Data Link connection released, and resumes idle state processing. When the remote SNDCF receives the corresponding DL-DISCONNECT indication, the Data Link Connection release phase is complete and the remote SNDCF resumes idle state processing as well.

Data Link Connection release may also be initiated by the DLS provider or by the called DLS user to refuse a connection. The action taken by the local SNDCEF in these circumstances is the same as described above. The mechanism and timing for releasing a Data Link Connection following the transmission of SN-Userdata by the SNDCEF are also local matters. Examples of circumstances which would cause the SNDCEF to release a Data Link Connection are:

- a) the expiration of a timeout period following the transmission of one or more PDUs;
- b) the need to use a specific interface to open an alternate Data Link Connection from the local Network entity to a different remote Network entity;
- c) the explicit intervention of the network management system; or
- d) a provider-initiated release of a Data Link Connection.

NOTE It is not a requirement that Data Link Connections be dynamically opened or closed for the correct operation of the SNDCEF herein described, the use of permanent Data Link Connections or the maintenance of Data Link Connections in a open state from system initialization is not precluded.


The selection of timeout values is a local matter. Factors to be considered in determining timeout values are the same as those identified in clause 8.4.3.4.

0 Introduction

ISO 8473 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. In particular, it is a protocol of the Network Layer. This International Standard may be used between Network entities in End systems or in Network Layer relay systems (or both). It provides the connectionless-mode Network Service as defined in ISO 8348/Add.1.

The interrelationship of these standards is illustrated in Figure 1.



1 Scope and field of application

This International Standard specifies a protocol which is used to provide the connectionless-mode Network-Service as described in ISO 8348/Add.1. The protocol relies upon the provision of an underlying connectionless-mode service by real subnetworks and/or data links.

The underlying connectionless-mode service assumed by the protocol may be obtained either directly, from a connectionless-mode real subnetwork, or indirectly, through the operation of an appropriate Subnetwork Dependent Convergence function (SND CF) or protocol (SND CP) over a connection-mode real subnetwork, as described in ISO 8648. This International Standard specifies the way in which this underlying connectionless-mode subnetwork service is obtained from real subnetworks that conform to ISO 8802-2 or ISO 8208. The underlying connectionless-mode subnetwork service may be obtained from real subnetworks that do not conform to either ISO 8802-2 or ISO 8208, but the way in which this is done is not specified by this International Standard.

NOTE The SND CF defined for the operation of this protocol in configurations that include a subnetwork using the X.25 PLP can be generalized in a straightforward way to cover the operation of ISO 8473 in configurations that include subnetworks that use connection-oriented protocols other than the X.25 PLP.

This International Standard specifies

- a) procedures for the connectionless transmission of data and control information from one Network entity to a peer Network entity;
- b) the encoding of the protocol data units (PDUs) used for the transmission of data and control information, comprising a variable-length protocol header format;
- c) procedures for the correct interpretation of protocol control information; and
- d) the functional requirements for implementations claiming conformance to this International Standard.

The procedures are defined in terms of

- a) the interactions among peer Network entities through the exchange of protocol data units;
- b) the interactions between a Network entity and a Network Service user through the exchange of Network Service primitives; and
- c) the interactions between a network-entity and an underlying service provider through the exchange of service primitives.

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 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/TR 8509, *Information processing systems — Open Systems Interconnection — Service conventions*.

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

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

ISO 8886, *Information processing systems — Data communication — Data link service definition for Open Systems Interconnection*¹⁾.

ISO 9074, *Information processing systems — Open systems interconnection — ESTELLE (Formal Description Technique based on an Extended State Transition Model)*¹⁾.

¹⁾ At present, at the stage of draft; publication anticipated in due course.

Section 1. General

3 Definitions

3.1 Reference Model definitions

This International Standard makes use of the following terms defined in ISO 7498:

- a) end-system
- b) network-entity
- c) Network layer
- d) Network protocol
- e) Network protocol data unit
- f) Network relay
- g) Network-service
- h) Network-service-access-point
- i) Network-service-access-point-address
- j) routing
- k) service
- l) service data unit
- m) service primitive

3.2 Service conventions definitions

This International Standard makes use of the following terms defined in ISO/TR 8509:

- a) service provider
- b) service user

3.3 Network Layer architecture definitions

This International Standard makes use of the following terms defined in ISO 8648:

- a) intermediate system
- b) relay system
- c) subnetwork
- d) subnetwork independent convergence protocol
- e) subnetwork independent convergence function
- f) subnetwork access protocol

3.4 Network Layer addressing definitions

This International Standard makes use of the following terms defined in ISO 8348/Add.2:

- a) Network addressing domain
- b) Network protocol address information
- c) subnetwork point of attachment

3.5 Local Area Network definitions

This International Standard makes use of the following terms defined in ISO 8802:

- a) local area network
- b) logical link control
- c) media access control

3.6 X.25 definitions

This International Standard makes use of the following terms defined in ISO 8208:

- a) data circuit-terminating equipment
- b) data terminal equipment
- c) logical channel
- d) permanent virtual circuit
- e) virtual circuit

3.7 Additional definitions

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

3.7.1

derived PDU

a protocol data unit whose fields are identical to those of an initial PDU, except that it carries only a segment of the user data from an N-UNITDATA request

3.7.2

initial PDU

a protocol data unit carrying the whole of the user data from an N-UNITDATA request

3.7.3

local matter

a decision made by a system concerning its behaviour in the Network Layer that is not prescribed or constrained by this International Standard

3.7.4

network-entity title

an identifier for a network-entity which has the same abstract syntax as an NSAP address, and which can be used to unambiguously identify a network-entity in an end or intermediate system

3.7.5

reassembly

the act of regenerating an initial PDU from two or more derived PDUs

3.7.6

segment

a distinct unit of data consisting of part or all of the user data provided in the N-UNITDATA request and delivered in the N-UNITDATA indication.

3.7.7

segmentation

the act of generating two or more derived PDUs from an initial or derived PDU. The derived PDUs together carry the entire user data of the initial or derived PDU from which they were generated

4 Abbreviations

4.1 Data Units

NSDU	Network service data unit
PDU	protocol data unit
SDU	service data unit
SNSDU	subnetwork service data unit

4.2 Protocol Data Units

DT PDU	data protocol data unit
ER PDU	error report protocol data unit

4.3 Protocol Data Unit Fields

CS	checksum
DA	destination address
DAL	destination address length
DUID	data unit identifier
E/R	error report flag
LI	length indicator
LT	lifetime
MS	more segments flag
NLPID	network layer protocol identifier
SA	source address
SAL	source address length
SL	segment length
SO	segment offset
SP	segmentation permitted flag
TL	total length
TP	type
V/P	version/protocol identifier Extension

4.4 Parameters

DA	destination address
QoS	quality of service
SA	source address

4.5 Miscellaneous

CLNP	connectionless-mode network protocol (i.e. the protocol defined in this International Standard)
DCE	data circuit-terminating equipment
DTE	data terminal equipment
LLC	logical link control
MAC	media access control
NS	network-service
NPAI	Network protocol address information
NSAP	Network service access point
PVC	permanent virtual circuit
SN	subnetwork
SNDCF	subnetwork dependent convergence function
SNDCP	subnetwork dependent convergence protocol
SNICP	subnetwork independent convergence protocol
SNAcP	subnetwork access protocol
SNPA	subnetwork point of attachment
SNCR	subnetwork connection reference
UI	unnumbered information

5 Overview of the protocol

5.1 Internal organization of the Network Layer

The architectural organization of the Network Layer is described in a separate International Standard, ISO 8648. ISO 8648 identifies and categorizes the way in which functions can be performed within the Network Layer by Network Layer protocols, thus providing a uniform framework for describing how protocols operating either individually or cooperatively in the Network Layer can be used to provide the OSI Network-Service. This protocol is designed to be used in the context of the internetworking protocol approach to the provision of the connectionless-mode Network-Service defined in ISO 8648.

This protocol is intended for use in the Subnetwork Independent Convergence protocol (SNICP) role. A protocol which fulfills the SNICP role operates to construct the OSI Network-Service over a defined set of underlying services, performing functions which are necessary to support the uniform appearance of the OSI connectionless-mode Network-Service over a homogeneous or heterogeneous set of interconnected subnetworks. This protocol is defined to accommodate variability where Subnetwork Dependent Convergence protocols and/or Subnetwork Access protocols do not provide all of the functions necessary to support the connectionless-mode Network-Service over all or part of the path from one NSAP to another.

As described in ISO 8648, a protocol at the Network Layer may fulfill different roles in different configurations. Although this protocol is designed particularly to be suitable for a SNICP role in the context of the internetworking protocol approach to the provision of the connectionless-mode Network-Service, it may also be used to fulfill other roles and may therefore be used in the context of other approaches to subnetwork inter-connection.

The operation of this protocol is specified with respect to an underlying service which is made available through the operation of other Network Layer protocols or through provision of the Data Link Service. The underlying service assumed by this protocol is described in 5.5.

5.2 Subsets of the protocol

Two subsets of the full protocol are defined which exploit the known subnetwork characteristics of particular configurations and are therefore not subnetwork independent.

The *Inactive Network Layer* protocol subset is a null-function subset which can be used when it is known that the source and destination end-systems are connected by a single subnetwork, and when none of the functions performed by the full protocol is required to provide the connectionless-mode Network-Service between any pair of end-systems.

The *Non-segmenting* protocol subset permits simplification of the header when it is known that the source and destination end-systems are connected by subnetworks whose individual service data unit sizes are greater than or equal to a known bound which is large enough so that segmentation is not required. This subset is selected by setting the *Segmentation Permitted* tag to zero (see 6.7).

5.3 Addresses and titles

The following subclauses describe the addresses and titles used by this protocol.

5.3.1 Addresses

The *Source Address* and *Destination Address* parameters referred to in 7.3 of this International Standard are OSI Network Service Access Point (NSAP) Addresses. The syntax and semantics of an OSI Network-Service-Access-Point-Address are described ISO 8348/Add.2.

The encoding used by this protocol to convey NSAP Addresses is the preferred binary encoding specified in ISO 8348/Add.2. The NSAP address, encoded as a string of binary octets according to ISO 8348/Add.2, is conveyed in its entirety in the address fields described in 7.3.

5.3.2 Network-entity titles

A network-entity title is an identifier for a network-entity in an end-system or intermediate system. Network-entity titles are allocated from the same name space as NSAP addresses, and the determination of whether a name is an NSAP address or a network-entity title depends on the context in which the name is interpreted. The values of the *Source Routeing* and *Recording of Route* parameters defined in 7.5.4 and 7.5.5 respectively are network-entity titles. The values of the *Source Address* and *Destination Address* parameters in the Error Report PDU defined in 7.9 are also network-entity titles.

The encoding used by this protocol to convey network-entity titles is the preferred binary encoding specified in ISO 8348/Add.2. The network-entity title, encoded as a string of binary octets according to ISO 8348/Add.2, is conveyed in its entirety in the fields described in 7.5.4, 7.5.5 and 7.9.1.2.

5.4 Service provided by the protocol

This protocol provides the connectionless-mode Network Service described in ISO 8348/Add.1. The Network-Service primitive provided is summarized in Table 1.

NOTE ISO 8348/Add.1 states that the maximum size of a connectionless-mode Network-service-data-unit (NSDU) is 64512 octets.

5.5 Underlying service assumed by the protocol

It is intended that ISO 8473 be capable of operating over connectionless-mode services derived from a wide variety of real subnetworks and data links. Therefore, in order to simplify the specification of the protocol, its operation is defined (in clause 6) with respect to an abstract “underlying service” rather than any particular real subnetwork service. This underlying service consists of a single

SN-UNITDATA primitive which conveys the source and destination subnetwork point of attachment addresses, a subnetwork Quality of Service parameter, and a minimum number of octets of user data.

The SN-UNITDATA primitive is 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.

The primitive provided is summarized in Table 2.

Provision of the underlying service is described in clause 8.

5.6 Service assumed from the local environment

A timer service shall be provided to allow the protocol entity to schedule events.

Table 1 — Connectionless-mode Network-Service primitive

Primitive	Parameters
N-UNITDATA .Request .Indication	NS-Source-Address, NS-Destination-Address, NS-Quality-of-Service, NS-Userdata

Table 2 — Underlying Service primitive

Primitive	Parameters
SN-UNITDATA .Request .Indication	SN-Source-Address, SN-Destination-Address, SN-Quality-of-Service, SN-Userdata

There are three primitives associated with the S-TIMER service:

- the S-TIMER Request;
- the S-TIMER Response; and
- the S-TIMER Cancel.

The primitives provided are summarized in Table 3.

Table 3 — Timer primitive

Primitive	Parameters
S-TIMER .Request	S-Time, S-Name, S-Subscript
.Response	S-Name, S-Subscript

The S-TIMER Request primitive indicates to the local environment that it should initiate a timer of the specified name and subscript and maintain it for the duration specified by the time parameter.

The S-TIMER Response primitive is initiated by the local environment to indicate that the amount of time requested by the corresponding S-TIMER Request primitive has elapsed.

The S-TIMER Cancel primitive is an indication to the local environment that the specified timer(s) should be canceled. If the S-Subscript parameter is not specified, then all timers with the specified name are canceled; otherwise, the timer of the given name and subscript is cancelled. If no timers correspond to the parameters specified, the local environment takes no action.

The parameters of the S-TIMER service primitives are specified in Table 3.

The S-Time parameter indicates the time duration of the specified timer. An identifying label is associated with a timer by means of the S-Name parameter. The S-Subscript parameter specifies a value to distinguish timers with the same name. The name and subscript taken together constitute a unique reference to the timer.

Timers used in association with a specific protocol function are defined under that protocol function.

NOTE This International Standard does not define specific values for the timers. Any derivations described in this Standard are not mandatory. Timer values should be chosen so that the requested Quality of Service can be provided, given the known characteristics of the underlying service.

Section 2. Specification of the protocol

6 Protocol functions

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

Not all of the functions must be performed by every implementation. Subclause 6.19 specifies which functions may be omitted, and the correct behaviour when requested functions are not implemented.

6.1 PDU Composition function

This function is responsible for the construction of a protocol data unit according to the rules governing the encoding of PDUs given in clause 7. Protocol Control Information required for delivering the data unit to its destination is determined from current state and local information and from the parameters associated with the N-UNITDATA Request.

Network Protocol Address Information (NPAI) for the *Source Address* and *Destination Address* fields of the PDU header is derived from the NS-Source-Address and NS-Destination-Address parameters.

The NS-Destination-Address and NS-Quality-of-Service parameters, together with current state and local information, are used to determine which optional functions are to be selected. User data passed from the Network-Service User (NS-Userdata) form the Data Part of the protocol data unit.

During the composition of the protocol data unit, a *Data Unit Identifier* is assigned to distinguish this request to transmit NS-Userdata to a particular destination NS User from other such requests. The originator of the PDU shall choose the Data Unit Identifier so that it remains unique (for this Source and Destination address pair) for the maximum lifetime of the Initial PDU in the network; this rule applies for any PDUs derived from the Initial PDU as a result of the application of the Segmentation function (see 6.7). Derived PDUs are considered to correspond to the same Initial PDU, and hence the same N-UNITDATA.Request, if they have the same Source Address, Destination Address, and Data Unit Identifier.

The Data Unit Identifier is also available for ancillary functions such as error reporting (see 6.10).

The total length of the PDU in octets is determined by the originator and placed in the *Total Length* field of the PDU header. This field is not changed in any Derived PDU for the lifetime of the protocol data unit.

When the non-segmenting protocol subset is employed, neither the Total Length field nor the Data Unit Identifier field is present. The rules governing the PDU composition function are modified in this case as follows. During the composition of the protocol data unit, the total length of the PDU in octets is determined by the originator and placed in the *Segment Length* field of the PDU header. This field is not changed for the lifetime of the PDU. No Data Unit Identification is provided.

6.2 PDU Decomposition function

This function is responsible for removing the Protocol Control Information from the protocol data unit. During this process, information pertinent to the generation of the N-UNITDATA.Indication is determined as follows. The NS-Source-Address and NS-Destination-Address parameters of the N-UNITDATA.Indication are recovered from the NPAI in the *Source Address* and *Destination Address* fields of the PDU header. The Data Part of the received PDU is retained until all segments of the original service data unit have been received; collectively, these form the NS-Userdata parameter of the N-UNITDATA.Indication. Information relating to the Quality of Service provided during the transmission of the PDU is determined from the Quality of Service and other information contained in the Options Part of the PDU header. This information constitutes the NS-Quality-of-Service parameter of the N-UNITDATA.Indication.

6.3 Header Format Analysis function

This function determines whether the full protocol described in this Standard is employed, or one of the defined proper subsets thereof. If the protocol data unit has a Network Layer Protocol Identifier indicating that this is a standard version of the protocol, this function determines whether a received PDU has reached its destination, using the Destination Address in the PDU header. If the Destination Address provided in the PDU identifies an NSAP served by this network-entity, then the PDU has reached its destination; if not, it shall be forwarded.

If the protocol data unit has a Network Layer Protocol Identifier indicating that the Inactive Network Layer protocol subset is in use, then no further analysis of the PDU header is required. The network-entity in this case determines that either the Subnetwork Point of Attachment address encoded as network protocol address information in the supporting subnetwork protocol corresponds directly to an NSAP address serviced by this network-entity or that an error has occurred.

6.4 PDU Lifetime Control function

This function is used to enforce the maximum PDU lifetime. It is closely associated with the Header Format Analysis function. This function determines whether a PDU received may be forwarded or whether its assigned lifetime has expired, in which case it shall be discarded.

The operation of the PDU Lifetime Control function depends upon the *Lifetime* field in the PDU header. This field contains, at any time, the remaining lifetime of the PDU (represented in units of 500 ms). The lifetime of the Initial PDU is determined by the originating network-entity and placed in the Lifetime field of the PDU. When the Segmentation function is applied to a PDU, the value of the Lifetime field of the Initial PDU is copied into all of the Derived PDUs.

The Lifetime of the PDU is decremented by every network-entity which processes the PDU. When a network-entity processes a PDU, it decrements the PDU Lifetime by at least one. The value of the PDU Lifetime field shall be decremented by more than one if the sum of

- a) the transit delay in the underlying service from which the PDU was received; and
- b) the delay within the system processing the PDU.

exceeds or is estimated to exceed 500 ms. In this case, the Lifetime field should be decremented by one for each additional 500 ms of delay. The determination of delay need not be precise, but where a precise value cannot be ascertained, the value used shall be an overestimate, not an underestimate.

If the Lifetime field reaches a value of zero before the PDU is delivered to the destination, the PDU shall be discarded. The Error Reporting function shall be invoked as described in 6.10. This may result in the generation of an Error Report PDU.

It is a local matter whether or not the destination network-entity performs the Lifetime Control function.

6.5 Route PDU function

This function determines the network-entity to which a protocol data unit should be forwarded and the underlying service that must be used to reach that network-entity, using the Destination Address and the Total Length fields of the PDU. Where segmentation is required, the Route PDU function further determines over which underlying service Derived PDUs/segments shall be sent in order to reach that network-entity. The results of the Route PDU function are passed to the Forward PDU function (along with the PDU itself) for further processing.

Selection of the underlying service that shall be used to reach the “next” system in the route is initially influenced by the NS-Quality-of-Service parameter of the N-UNITDATA.Request, which specifies the QoS requested by the sending NS User. Whether this QoS is to be provided directly by the CLNP, through the selection of the N-Quality-of-Service parameter and other optional parameters, or through the QoS facilities offered by each of the underlying services is determined prior to invocation of the Forward PDU function. Route selection by intermediate systems may subsequently be influenced by the values of the N-Quality-of-Service parameter (if present), and other optional parameters (if present).

6.6 Forward PDU function

This function issues an SN-UNITDATA.Request primitive (see 5.5), supplying the subnetwork or SNDCF identified by the Route PDU function with the protocol data unit as user data to be transmitted, the address information required by that subnetwork or SNDCF to identify the “next” system within the subnetwork-specific addressing domain (this may be an intermediate system or the destination end-system), and Quality of Service constraints (if any) to be considered in the processing of the user data.

When the PDU to be forwarded is longer than the maximum service data unit size provided by the underlying service, the Segmentation function is applied (see 6.7).

6.7 Segmentation function

Segmentation is performed when the size of the protocol data unit is greater than the maximum service data unit size supported by the underlying service to be used to transmit the PDU.

Segmentation consists of composing two or more new PDUs (Derived PDUs) from the received PDU. The received PDU may be the Initial PDU, or it may be a Derived PDU. All of the header information from the PDU to be segmented, with the exception of the segment length and checksum fields of the fixed part, and the segment offset of the segmentation part, is duplicated in each Derived PDU, including all of the address part, the data unit identifier and total length of the segmentation part, and the options part (if present).

NOTE The rules for forwarding and segmentation guarantee that the header length is the same for all segments (Derived PDUs) of the Initial PDU, and is the same as the header length of the Initial PDU. The size of a PDU header will not change due to the operation of any protocol function.

The user data encapsulated within the received PDU are divided such that the Derived PDUs satisfy the size requirements of the SN-Userdata parameter field of the SN-UNITDATA.Request primitive used to access the underlying service. Each derived PDU, except for the last, shall contain a number of octets that is a non-zero multiple of 8. Thus, the value of the segment offset field in any PDU is either zero or a non-zero multiple of 8. Segmentation shall not result in the generation of a Derived PDU containing less than eight octets of user data.

Derived PDUs are identified as being from the same Initial PDU by means of

- a) the Source Address field;
- b) the Destination Address field; and
- c) the Data Unit Identifier field.

The following fields of the PDU header are used in conjunction with the Segmentation function:

- a) *Segment Offset* — identifies the octet at which the segment begins with respect to the start of the Initial PDU;
- b) *Segment Length* — specifies the number of octets in the Derived PDU, including both header and data;
- c) *More Segments Flag* — is set to one if this Derived PDU does not contain the final octet of the Initial PDU as its final octet of user data; and
- d) *Total Length* — specifies the entire length of the Initial PDU, including both header and data.

Derived PDUs may be further segmented without constraining the routing of the individual Derived PDUs.

The *Segmentation Permitted* flag is set to one to indicate that segmentation is permitted. If the Initial PDU is not to be segmented at any point during its lifetime in the network, the flag is set to zero by the source network-entity. The setting of the Segmentation Permitted flag may not be changed by any other network-entity for the lifetime of the Initial PDU and any Derived PDUs.

6.8 Reassembly function

The Reassembly function reconstructs the Initial PDU from the Derived PDUs generated by the operation of the Segmentation function on the Initial PDU (and, recursively, on subsequent Derived PDUs).

A bound on the time during which segments (Derived PDUs) of an Initial PDU will be held at a reassembly point before being discarded is provided, so that reassembly resources may be released when it is no longer expected that any outstanding segments of the Initial PDU will arrive at the reassembly point. Upon reception of a Derived PDU, a reassembly timer is initiated with a value which indicates the amount of time which shall elapse before any outstanding segments of the Initial PDU shall be assumed to be lost. When this timer expires, all segments (Derived PDUs) of the Initial PDU held at the reassembly point are discarded, the resources allocated for those segments are freed and, if selected, an Error Report is generated (see 6.10).

While the exact relationship between reassembly lifetime and PDU lifetime is a local matter, the Reassembly function shall preserve the intent of the PDU lifetime. Consequently, the reassembly function shall discard PDUs whose lifetime would otherwise have expired had they not been under the control of the reassembly function; that is, the reassembly lifetime for a given PDU shall be less than the PDU lifetime in all derived PDUs being held at the reassembly point.

NOTE 1 Methods of bounding reassembly lifetime are discussed in Annex B.

NOTE 2 The Segmentation and Reassembly functions are intended to be used in such a way that the fewest possible segments are generated at each segmentation point and reassembly takes place at the final destination of a PDU. However, other schemes which

- a) interact with the routing algorithm to favour paths on which fewer segments are generated;
- b) generate more segments than absolutely required in order to avoid additional segmentation at some subsequent point; or
- c) allow partial or full reassembly at some intermediate point along the route

are not precluded. The information necessary to enable the use of one of these alternative strategies may be made available through the operation of a Network Layer Management function or by other means.

NOTE 3 The originator of the Initial PDU determines the value of the *Segmentation Permitted* flag in the Initial PDU and all Derived PDUs (if any). Partial or full reassembly in an intermediate system [NOTE 2 c) above] may not change this value in the Initial PDU or any PDU derived from it, and may not therefore add or remove the segmentation part of the header.

6.9 Discard PDU function

This function performs all of the actions necessary to free the resources reserved by the network-entity when any of the following situations is encountered (It should be noted that the list is not exhaustive):

- a) A violation of protocol procedure has occurred.
- b) A PDU is received whose checksum is inconsistent with its contents.
- c) A PDU is received, but due to local congestion, it cannot be processed.
- d) A PDU is received whose header cannot be analyzed.
- e) A PDU is received which cannot be segmented and cannot be forwarded because its length exceeds the maximum service data unit size supported by any underlying service available for transmission of the PDU to the next network-entity on the chosen route.
- f) A PDU is received whose destination address is unreachable or unknown.
- g) Incorrect or invalid source routing was specified. This may include a syntax error in the source routing field, an unknown or unreachable network-entity title in the source routing field, or a path which is not acceptable for other reasons.
- h) A PDU is received whose PDU lifetime has expired or whose lifetime expires during reassembly.
- i) A PDU is received which contains an unsupported option.

6.10 Error Reporting function

6.10.1 Overview

This function attempts to return an Error Report PDU to the source network-entity when a protocol data unit is discarded in accordance with 6.9.

The Error Report PDU identifies the discarded PDU, specifies the type of error detected, and identifies the location in the header of the discarded PDU at which the error was detected. At least the entire header of the Discarded PDU and, at the discretion of the originator of the Error Report PDU, none, all, or part of the Data Part is placed in the Data Part of the Error Report PDU.

The originator of a Data PDU controls the generation of Error Report PDUs. An *Error Report* flag in the original PDU is set by the source network-entity to indicate that an Error Report PDU is to be returned if the Initial PDU or any PDUs derived from it are discarded; if the flag is not set, Error Reports are not generated.

NOTE 1 The suppression of Error Report PDUs is controlled by the originating network-entity and not by the NS User. Care should be exercised by the originator with regard to suppressing ER PDUs so that error reporting is not suppressed for every PDU generated.

NOTE 2 Non-receipt of an Error Report PDU does not imply correct delivery of a PDU issued by a source network-entity.