

*(Including addenda 1, 2
and 3)*

Information processing systems : Data communications : Network service definition

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 Division)
 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)
 Information Technology Users Standards Association
 Inter-universities Computing Committee
 Joint Network Team
 LAMSAC
 National Computing Centre Ltd.
 National Health Services
 OFTEL (Office of Telecommunications)
 Post Office
 Telecommunication Engineering and Manufacturing Association
 Telecommunications Managers' Association
 User Standards Forum for Information Technology (Institute of Data Processing Management)

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 November 1990

© BSI 1990

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

ISBN 0 580 19099 4

Amendments issued since publication

Amd. No.	Date	
7189	7.7.92	Id. Hicks

National foreword

This British Standard, prepared under the direction of the Information Systems Technology Standards Policy Committee, is identical with ISO 8348 : 1987 'Information processing systems — Data communications — Network service definition' including Technical Corrigendum 1 published by the International Organization for Standardization (ISO). It supersedes DD 118, DD 119, and DD 134 which are withdrawn.



Cross-references

International standard	Corresponding British Standard
ISO 7498 : 1984	BS 6568 : Reference model of Open System Interconnection Part 1 : 1988 Basic reference model (incorporating connectionless-mode transmission) (Identical)
ISO/TR 8509 : 1987	DD 108 : 1988 Open Systems Interconnection Service conventions (Identical)

The Technical Committee has reviewed the provisions of ISO 8073 and CCITT Recommendations X.213, to which reference is made in the text, and has decided that they are acceptable for use in connection with this standard. A related standard to ISO 8073 is DD 116 'Open Systems Interconnection : transport protocol specification'. CCITT publications are available from BSI Sales Department, Linford Wood, Milton Keynes, MK14 6LE.

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

Contents

	Page
Committees responsible	Inside front cover
National foreword	ii
Method	
0 Introduction	1
1 Scope and field of application	1
2 References	2
Section one: General	
3 Definitions	3
4 Abbreviations	3
5 Conventions	3
6 Overview and general characteristics	4
7 Features of the Network Service	4
8 Classes of Network Service	4
9 Model of the Network Service	5
10 Quality of Network Service	8
Section two: Definition of the connection-mode primitives	
11 Sequence of primitives	12
12 Network Connection establishment phase	15
13 Network Connection release phase	19
14 Data transfer phase	21
Annex - Differences between ISO 8348 and CCITT Recommendation X.213-1986	25
Addenda	
1 Connectionless-mode transmission	27
2 Network layer addressing	41
3 Additional features of the network service	61

Information processing systems — Data communications — Network service definition

0 Introduction

This International Standard is one of a set of International Standards produced to facilitate the interconnection of computer systems. It is related to other International Standards in the set as defined by the Open Systems Interconnection (OSI) Basic Reference Model. The OSI Reference Model subdivides the area of standardization for interconnection into a series of layers of specification, each of a manageable size.

This International Standard defines the Service provided by the Network Layer to the Transport Layer at the boundary between the Network and Transport Layers of the Reference Model. It provides for the designers of Transport Protocols a definition of the Network Service existing to support the Transport Protocol and for the designers of Network Protocols a definition of the services to be made available through the action of the Network Protocol over the underlying service. This relationship is illustrated in figure 1.




Figure 1 — Relationship of the Network Service in this International Standard to the protocols specified in other OSI Standards

The use of the word "Network" to name the "Network" Layer of the OSI Reference Model should be distinguished from the use of the word "network" to denote a communications network as conventionally understood. To facilitate this distinction, the term "subnetwork" is used for a collection of physical equipment, commonly called a "network" (see ISO 7498). Subnetworks may be either public networks or privately supplied networks. In the case of public networks, their properties may be determined by separate CCITT Recommendations such as CCITT Recommendation X.21 for a circuit-switched network or CCITT Recommendation X.25 for a packet-switched network.

Throughout the set of OSI Standards the term "Service" refers to the abstract capability provided by one layer of the OSI Reference Model to the layer above it. Thus, the Network Service defined in this International Standard is a conceptual architectural Service, independent of administrative divisions.

NOTE — It is important to distinguish the specialized use of the term "Service" within the set of OSI Standards from its use elsewhere to describe the provision of a service by an organization (such as the provision of a service, as defined in other CCITT Recommendations, by an Administration).

Any particular subnetwork may or may not support the OSI Network Service. The OSI Network Service may be provided by a combination of one or more subnetworks and optional additional functions between or outside these subnetworks.

1 Scope and field of application

This International Standard defines the OSI Network Service in terms of

- the primitive actions and events of the Service;
- the parameters associated with each primitive action and event, and the form which they take;
- the interrelationship between, and the valid sequences of, these actions and events.

The principal objectives of this International Standard are

- to specify the characteristics of a conceptual Network Service and thus supplement the Reference Model in guiding the development of Network Layer protocols;
- to encourage convergence of the capabilities offered by providers of subnetworks;
- to provide a basis for the individual enhancement of existing heterogeneous subnetworks to a common subnetwork-independent Network Service to enable them to be concatenated for the purpose of providing global communication. (Such concatenation may involve optional additional functions which are not defined in this International Standard.) A definition of the quality of service is an important element of this International Standard;

ISO 8348 : 1987 (E)

d) to provide a basis for the development and implementation of subnetwork-independent Transport Layer protocols decoupled from the variability of underlying public and private subnetworks and their specific interface requirements.

This International Standard does not specify individual implementations or products nor does it constrain the implementation of entities and interfaces within a system.

There is no conformance of equipment to this International Standard. Instead, conformance is achieved through implementation of conforming OSI Network protocols which fulfill the Network Service defined in this International Standard.

2 References

ISO 7498, *Information processing systems — Open Systems Interconnection — Basic Reference Model*.

NOTE — See also CCITT Recommendation X.200, *Reference Model of Open Systems Interconnection for CCITT Applications*.

ISO 8073, *Information processing systems — Open Systems Interconnection — Connection oriented transport protocol specification*.

NOTE — See also CCITT Recommendation X.224, *Transport Protocol Specification for Open Systems Interconnection for CCITT Applications*.

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

NOTE — See also CCITT Recommendation X.210, *OSI Layer Service definition conventions*.

CCITT Recommendation X.213, *Network Service Definition for Open Systems Interconnection for CCITT Applications*.

Section one : General

3 Definitions

NOTE — The definitions contained in this clause make use of abbreviations defined in clause 4.

3.1 Reference Model definitions

This International Standard is based on the concepts developed in ISO 7498, and makes use of the following terms defined in it:

- a) expedited Network-Service-data-unit
- b) Network-Connection
- c) Network Layer
- d) Network Service
- e) Network-Service-access-point
- f) Network-Service-access-point-address
- g) Network-Service-data-unit
- h) subnetwork

3.2 Service conventions definitions

This International Standard also makes use of the following terms defined in ISO/TR 8509, as they apply to the Network Layer:

- a) Network Service user
- b) Network Service provider
- c) primitive
- d) request
- e) indication
- f) response
- g) confirm

3.3 Network Service definitions

For the purpose of this International Standard, the following definitions also apply.

3.3.1 Calling NS user: An NS user that initiates an NC Establishment request.

3.3.2 Called NS user: An NS user with whom a calling NS user wishes to establish an NC.

NOTE — Calling NS users and called NS users are defined with respect to a single NC. An NS user can be both a calling and a called NS user simultaneously.

3.3.3 Generic address: An address which identifies a set of NSAPs rather than a single specific NSAP.

4 Abbreviations

COR	Confirmation of Receipt
ENSDU	Expedited Network-Service-Data-Unit
N	Network
NC	Network-connection
NL	Network Layer
NS	Network Service
NSAP	Network-Service-access-point
NSDU	Network-Service-data-unit
OSI	Open Systems Interconnection
QOS	Quality of service

5 Conventions

5.1 General conventions

This International Standard uses the descriptive conventions defined in ISO/TR 8509.

The layer service model, service primitives, and time-sequence diagrams taken from those conventions are entirely abstract descriptions; they do not represent a specification for implementation.

5.2 Parameters

Service primitives, used to represent service-user/service-provider interactions (see ISO/TR 8509), convey parameters which indicate information available in the user/provider interaction.

The parameters which apply to each group of Network Service primitives are set out in tables in clauses 12 to 14. Each "X" in the tables indicates that the primitive labelling the column in which it falls may carry the parameter labelling the row in which it falls.

Some entries are further qualified by items in brackets. These may be

- a) an indication that the parameter is conditional in some way;
- (C) indicates that the parameter is not present on the primitive for every NC; the parameter definition describes the conditions under which the parameter is present or absent.
- b) a parameter specific constraint:
- (=) indicates that the value supplied in an indication or confirm primitive is always identical to that supplied in the corresponding request or response primitive occurring at the peer NSAP.

ISO 8348 : 1987 (E)

- c) indication that some note applies to the entry:

(Note x) indicates that the referenced note contains additional information pertaining to the parameter and its use.

In any particular interface, not all parameters need be explicitly stated. Some may be implicitly associated with the NSAP at which the primitive is issued.

5.3 NC endpoint identification convention

If an NS user needs to distinguish among several NCs at the same NSAP, then a local NC endpoint identification mechanism must be provided. All primitives issued at such an NSAP would be required to use this mechanism to identify NCs. Such an implicit identification is not described as a parameter of the service primitives in this International Standard.

NOTE — The implicit NC endpoint identification must not be confused with the address parameters of the N-CONNECT primitives (see 12.2).

6 Overview and general characteristics

The Network Service provides for the transparent transfer of data (i.e. NS-user-data) between NS users. It makes invisible to these NS users the way in which supporting communications resources are utilized to achieve this transfer.

In particular, the Network Service provides for the following:

- a) Independence of underlying transmission media — The Network Service relieves NS users from all concerns regarding how various subnetworks are used to provide the Network Service. The Network Service hides from the NS user differences in the transfer of data over heterogeneous subnetworks, other than quality of service.
- b) End-to-end transfer — The Network Service provides for transfer of NS-user-data between NS users in end systems. All routing and relaying functions are performed by the NS provider including the case where several similar or dissimilar transmission resources are used in tandem or in parallel.
- c) Transparency of transferred information — The Network Service provides for the transparent transfer of octet-aligned NS-user-data and/or control information. It does not restrict the content, format or coding of the information, nor does it ever need to interpret its structure or meaning.
- d) Quality of service selection — The Network Service makes available to NS users a means to request and to agree to the quality of service for the transfer of NS-user-data. Quality of service is specified by means of QOS-parameters representing characteristics such as throughput, transit delay, accuracy, and reliability.
- e) NS-user-addressing — The Network Service utilizes a system of addressing (NSAP addressing) which allows NS users to refer unambiguously to one another.

7 Features of the Network Service

The Network Service offers the following features to an NS user:

- a) The means to establish an NC with another NS user for the purpose of transferring NS-user-data in the form of NSDUs. More than one NC may exist between the same pair of NS users.
- b) The establishment of an agreement between the two NS users and the NS provider for a certain QOS associated with each NC.
- c) The means of transferring NSDUs in sequence on an NC. The transfer of NSDUs, which consist of an integer number of octets, is transparent, in that the boundaries of NSDUs and the contents of NSDUs are preserved unchanged by the Network Service, and there are no constraints on the NSDU content imposed by the Network Service.
- d) The means by which the receiving NS user may flow control the rate at which the sending NS user may send NSDUs.
- e) In some circumstances, the means of transferring separate expedited NSDUs in sequence (see clause 8). Expedited NSDUs are limited in length and their transmission is subject to a different flow control from normal data across the NSAP.
- f) The means by which the NC can be returned to a defined state and the activities of the two NS users synchronized by use of a Reset service.
- g) In some circumstances, the means for the NS user to confirm the receipt of an NSDU (see clause 8).
- h) The unconditional, and therefore possibly destructive, release of an NC by either of the NS users or by the NS provider.

8 Classes of Network Service

No distinct classes of Network Service are defined. However, two Network Layer services, Receipt Confirmation and Expedited Data Transfer, are NS provider-options.

A service which is an NS provider-option is one which an NS provider can choose either to provide or not to provide for a particular NC. In circumstances where the NS provider chooses not to provide a provider-option service, it will not be available in the Network Service. If the provider-option Receipt Confirmation or Expedited Data Transfer is provided, it shall be provided as defined in this International Standard (see clauses 14.1 to 14.3).

All other Network services are mandatory in the Network Service. Mandatory services shall be provided by every NS provider and are therefore always available.

9 Model of the Network Service

9.1 Model of the Network Layer Service

This International Standard uses the abstract model for a layer service defined in clause 4 of ISO/TR 8509. The model defines the interactions between the NS users and the NS provider which take place at the two NSAPs. Information is passed between the NS user and the NS provider by service primitives, which may convey parameters.

There are two types of OSI Network Service:

- a) a connection-mode Service (defined in section two). The connection-mode Service is characterized by the features a) to h) given in clause 7.
- b) a connectionless-mode Service (defined in addendum 1 to this International Standard).

When making reference to the Network Service, an NS user or NS provider shall state which types of Network Service it expects to use or provide.

9.2 Model of a Network Connection

Between the two endpoints of an NC, there exists a flow control function which relates the behaviour of the NS user at one end receiving NS-user-data to the ability of the NS user at the other end to send NS-user-data. As a means of specifying this flow control feature and its relationship with other capabilities provided by the Network Service, the queue model of an NC, described in the following clauses, is used.

This queue model of an NC is discussed only to aid in the understanding of the end-to-end service features perceived by users of the Network Service. It is not intended to serve as a substitute for a precise, formal description of the Network Service, nor as a complete specification of all allowable sequences of NS primitives. (Allowable primitive sequences are specified in clause 11 — also, see the note.) In addition, this model does not attempt to describe all the functions or operations of Network Layer entities (including relay entities) which are used to

provide the Network Service. No attempt to specify or constrain Network Service implementations is implied.

In interpreting this International Standard, statements in clauses 12 to 14 concerning the properties of individual primitives have precedence over the general statements in this clause.

NOTE — In addition to the interaction between service primitives described by this model, there may be constraints applied locally on the ability to invoke primitives, as well as service procedures defining particular sequencing constraints on some primitives.

9.2.1 Queue model concepts

The queue model represents the operation of an NC in the abstract by a pair of queues linking the two NSAPs. There is one queue for each direction of information flow (see figure 2).

Each queue represents a flow control function in one direction of transfer. The ability of an NS user to add objects to a queue will be determined by the behaviour of the NS user removing objects from that queue and the state of the queue. Objects are entered or removed from the queue, either as the result of interactions at the two NSAPs, or as the result of NS provider initiatives.

The pair of queues is considered to be available for each potential NC.

The objects which may be placed in a queue as a result of interactions at an NSAP (see clauses 12 to 14) are

- a) connect objects (associated with N-CONNECT primitives and all of their parameters);
- b) octets of normal NS-user-data (associated with an N-DATA primitive);
- c) indications of end-of-NSDU (associated with completion of an N-DATA primitive);
- d) expedited NSDUs (associated with N-EXPEDITED-DATA primitives and all their parameters);
- e) data acknowledgment objects (associated with N-DATA-ACKNOWLEDGE primitives);




Figure 2 — Queue model of a Network Connection

- f) reset objects (associated with N-RESET primitives and their parameters);
- g) disconnect objects (associated with N-DISCONNECT primitives and all their parameters).

NOTE — The description of flow control (see 9.2.3) requires a less abstract description than that used for describing sequences of primitives in clauses 11 to 14. While primitives are defined to be indivisible, for purposes of this queue model, information associated with N-DATA primitives is conceptually subdivided into a sequence of octets of NS-user-data followed by an end-of-NSDU indication. This does not imply any particular subdivision in any real interface.

The objects which may be placed in a queue as a result of NS provider initiatives (see clauses 12 to 14) are

- 1) reset objects (associated with N-RESET primitives and all their parameters);
- 2) synchronization mark objects (see 9.2.4);
- 3) disconnect objects (associated with N-DISCONNECT primitives and all their parameters).

The queues are defined to have the following general properties:

- i) A queue is empty until a connect object has been entered and can be returned to this state, with loss of its contents, by the NS provider (see 9.2.4 and 9.2.5).
- ii) Objects may be entered into a queue as a result of the actions of the source NS user, subject to control by the NS provider; objects may also be entered into a queue by the NS provider.
- iii) Objects are removed from the queue under the control of the receiving NS user.
- iv) Objects are normally removed under the control of the NS user in the same order that they were entered (however see 9.2.3).
- v) A queue has a limited capacity, but this capacity is not necessarily either fixed or determinable.

9.2.2 NC establishment

A pair of queues is associated with an NC between two NSAPs when the NS provider receives an N-CONNECT request primitive at one of the NSAPs and a connect object is entered into one of the queues. From the standpoint of one of the NS users of the NC, the queues remain associated with the NC until a disconnect object (associated with an N-DISCONNECT primitive) is either entered or removed from a queue at that NSAP.

If NS user A denotes the NS user who initiates NC Establishment (resulting in a connect object being entered into the queue from NS user A to NS user B) then no object other than a disconnect object may be entered into the queue from A to B until after the connect object associated with the N-CONNECT confirm has been removed. In the queue from NS user B to NS user A, objects can be entered only after a connect object associated with an N-CONNECT response from NS user B has been entered; it is possible for a disconnect object to be placed in the queue from B to A, instead of a connect object, to release the NC.

The properties exhibited by the queues while the NC exists represent the agreements reached among the NS users and the NS provider during the NC Establishment procedure concerning quality of service and the use of the Receipt Confirmation and Expedited Data Transfer services.

9.2.3 Data transfer operations

Flow control on the NC is represented in this queue model by the management of the queue capacity, allowing objects of certain types to be added to the queues. The conditions affecting entry of reset and disconnect objects are described in item b) below and in 9.2.4 and 9.2.5. The flow control relationship between the other types of objects is summarized by table 1.

Once in the queue, the NS provider may manipulate pairs of adjacent objects, resulting in

- a) change of order; the order of any pair of objects may be reversed if, and only if, the following object is of a type defined to be able to advance ahead of the preceding object. No object is defined to be able to advance ahead of another object of the same type.

Table 1 — Flow control relationships between queue model objects

may prevent further addition of object y \ The addition of object x	Octets of NS-user-data or end-of-NSDU	Expedited NSDU	Data acknowledgement
Octets of normal NS-user-data or end-of-NSDU	Yes	Yes	No
Expedited NSDU	No	Yes	No
Data acknowledgement	No	No	No

b) deletion; any object may be deleted if, and only if, the following object is defined to be destructive with respect to the preceding object. If necessary, the last object in the queue will be deleted to allow a destructive object to be entered. Destructive objects may therefore always be added to the queue. Disconnect objects are defined to be destructive with respect to all other objects. Reset objects are defined to be destructive with respect to all other objects except connect, disconnect, and other reset objects.

The relationships between objects which may be manipulated as described in a) and b) are summarized in table 2.

Whether the NS provider performs actions resulting in change of order and deletion or not will depend upon the behaviour of the NS users and the agreed QOS for the NC. In general, if an NS user does not cause objects to be removed from a queue, the NS provider shall, after some unspecified period of time, perform all permitted actions of types a) and b).

9.2.4 Reset operations

The invocation of a Reset procedure is represented in the two queues as follows:

- a) invocation of a Reset procedure by the NS provider is represented by the introduction into each queue of a reset object followed by a synchronization mark object.
- b) A Reset procedure invoked by an NS user is represented by the addition of a reset object to one queue. In this case, the NS provider will insert a reset object followed by a synchronization mark object into the other queue.

The completion of a Reset procedure by the issuance of an N-RESET response by an NS user results in a reset object being placed in the queue from the responding NS user.

A synchronization mark object cannot be removed from a queue by an NS user; a queue appears empty to an NS user when a synchronization mark object is the next object in it. Unless destroyed by a disconnect object, a synchronization mark object remains in the queue until the next object following it in the queue is a reset object. Both the synchronization mark object and the following reset object are then deleted by the NS provider.

NOTE — Associated with the invocation of a Reset procedure are restrictions on the issuance of certain other types of primitives. These restrictions will result in restrictions on the entry of certain object types into the queue until the Reset procedure is complete.

9.2.5 NC Release

The insertion into a queue of a disconnect object, which may occur at any time, represents the initiation of an NC Release procedure. The Release procedure may be destructive with respect to other objects in the two queues and eventually results in the emptying of the queues and the disassociation of the queues with the NC.

The insertion of a disconnect object may also represent the rejection of an NC Establishment attempt or the failure to complete NC Establishment. In such cases, if a connect object representing an N-CONNECT request primitive is deleted by a disconnect object, then the disconnect object is also deleted. The disconnect object is not deleted when it deletes any other object, including the case where it deletes a connect object representing an N-CONNECT response.

Table 2 — Ordering relationships between queue model objects

with respect to preceding object y \ Following object x is defined	Connect	Octets of normal NS-user-data	End-of-NSDU	Expedited NSDU	Data acknowledgement	Reset	Synchronization mark	Disconnect
Connect	N/A	—	N/A	—	—	—	N/A	DES
Octets of normal NS-user-data	N/A	—	—	AA	AA	DES	N/A	DES
End of NSDU	N/A	—	N/A	AA	AA	DES	N/A	DES
Expedited NSDU	N/A	—	—	—	AA	DES	N/A	DES
Data acknowledgement	N/A	—	—	AA	—	DES	N/A	DES
Reset	N/A	—	N/A	—	—	DES	—	DES
Synchronization mark	N/A	—	—	—	—	DES	N/A	DES
Disconnect	N/A	N/A	N/A	N/A	N/A	N/A	N/A	DES

AA indicates that object x is defined to be able to advance ahead of the preceding object y.

DES indicates that object x is defined to be destructive with respect to the preceding object y.

— indicates that object x is neither destructive with respect to object y nor able to advance ahead of object y.

N/A indicates that object x will not occur in a position succeeding object y in a valid state of a queue.

10 Quality of Network Service

The term Quality of Service (QOS) refers to certain characteristics of an NC as observed between the NC end-points. QOS describes aspects of an NC which are attributable solely to the NS provider; it can only be properly determined in the absence of NS user behaviour (which is beyond the control of the NS provider) which specifically constrains or impairs the performance of the Network Service.

A value of QOS applies to an entire NC. When determined or measured at both ends of an NC, the QOS observed by the NS users at the two ends of the NC is the same. This is true even in the case of an NC spanning several subnetworks where each subnetwork offers different services.

10.1 Determination of QOS

QOS is described by means of QOS-parameters. The definition of each of these QOS-parameters specifies the way in which the QOS-parameter's value is measured or determined, making reference, where appropriate, to primitive events of the NS.

NOTES

- 1 It is important to distinguish the use of the term QOS-parameters from the more general term "parameters" as defined in 5.2 and used throughout this International Standard. A "QOS-parameter" refers to a specific aspect or component of the QOS for an NC. As described below, a particular QOS-parameter may or may not be related to a parameter defined as part of a Network Service primitive.
- 2 For purposes of accuracy and/or convenience, the definition and measurement formula for some QOS-parameters includes a component attributable to the NS user(s). In such cases, to evaluate the QOS attributable solely to the NS provider, this NS user-dependent component must be factored out.
- 3 The definition of NS' QOS-parameters in terms which provide a means for measurement should not be understood to imply that QOS monitoring or that verification of stated QOS values is, or must be, performed by the NS provider or by the NS users.

It is in terms of the NS QOS-parameters that information about QOS is exchanged among the NS provider and NS users.

Information about the QOS requirements of the NS users may be used by the NS provider for purposes such as protocol selection, route determination, and allocation of resources. Information about the QOS available from the NS provider may be used by NS users for purposes such as selecting QOS enhancement mechanisms and determining the QOS values provided to NS users at higher layers.

The NS QOS-parameters can be divided into two categories as follows:

- a) those whose values are "conveyed" between peer NS users by means of the NS during the establishment phase of an NC. As part of this conveyance, a three-party "negotiation" among the NS users and the NS provider for the purpose of agreeing upon a particular QOS-parameter value may take place; and
- b) those whose values are not "conveyed" or "negotiated" among the NS users and the NS provider. For

these QOS-parameters, however, information about the values which is useful to the NS provider and each NS user may be made known by local means.

The NS QOS-parameters are defined in 10.2.1 to 10.2.12.

The set of NS QOS-parameters that belong to the first category, and the procedures and constraints that apply to conveying and negotiating those QOS-parameters, are specified in 12.2.7. Once the NC is established, and throughout the lifetime of the NC, the agreed values for these two QOS-parameters are not "renegotiated" at any point, and there is no guarantee that the originally negotiated values will be maintained. The NS user should also be aware that, once an NC is established, changes in QOS on the NC are not explicitly signalled in the NS.

For QOS-parameters in the second category, the values for a particular NC are not negotiated, nor are they directly conveyed from NS user to NS user. As a local matter, however, there may be means by which the values of one or more of these QOS-parameters are known and utilized by the NS provider and each NS user. Despite the local nature of particular NS user/NS provider interactions which may occur for the purposes of exchanging QOS-parameter information, the characteristics of an NC which the QOS-parameters describe are applicable and can be observed on a complete NC, end-to-end basis. Thus, in order to give a full characterization of the properties of NCs, the definitions of the entire set of QOS-parameters which apply to the NS, including those classified in category 2, are included in this Network Service Definition. Other aspects related to category 2 parameters, such as the circumstances of their availability and use, as well as other QOS issues, such as the relationship to OSI management, and multi-layer QOS relationships, are the subjects of other OSI QOS-related specifications.

NOTE — For non-negotiated QOS-parameters associated with the Data Transfer phase of an NC, when specified, a value of such a QOS-parameter applies to both directions of transfer on the NC.

10.2 Definition of QOS-parameters

QOS-parameters can be classified as

- a) QOS-parameters which express Network Service performance, as shown in table 3.

Table 3 — Classification of performance QOS-parameters

Phase	Performance criterion	
	Speed	Accuracy/reliability
NC establishment	NC establishment delay	NC establishment failure probability (misconnection/NC refusal)
Data transfer	Throughput	Residual error rate (corruption, duplication/loss) NC resilience
	Transit delay	Transfer failure probability
NC release	NC release delay	NC release failure probability

- b) QOS-parameters that express other Network Service characteristics, as shown in table 4.

Table 4 — QOS-parameters not associated with performance

NC protection
NC priority
Maximum acceptable cost

NOTE — Some QOS-parameters are defined in terms of the issuance of Network Service primitives. Reference to a primitive in 10.2.1 to 10.2.12 refers to the complete execution of that service primitive at the appropriate NSAP.

10.2.1 NC establishment delay

NC establishment delay is the maximum acceptable delay between an N-CONNECT request and the corresponding N-CONNECT confirm primitive.

NOTE — This delay includes a component, attributable to the called NS user, which is the time between the N-CONNECT indication primitive and the N-CONNECT response.

10.2.2 NC establishment failure probability

NC establishment failure probability is the ratio of total NC establishment failures to total NC establishment attempts in a measurement sample.

NC establishment failure is defined to occur when a requested NC is not established within the specified maximum acceptable time period as a result of NS provider behaviour such as misconnection, NC refusal, or excessive delay. NC establishment attempts which fail as a result of NS user behaviour such as error, NC refusal, or excessive delay are excluded in calculating NC establishment failure probability.

10.2.3 Throughput

Throughput is defined, for each direction of transfer, in terms of a sequence of at least two successfully transferred NSDUs presented continuously to the NS provider at the maximum rate the NS provider can continuously sustain, and unconstrained by flow control applied by the receiving NS user.

Given such a sequence of n NSDUs, where n is greater than or equal to 2, the Throughput is defined to be the smaller of

- the number of NS-user-data octets contained in the last $n - 1$ NSDUs divided by the time between the first and last N-DATA requests in the sequence; and
- the number of NS-user-data octets contained in the last $n - 1$ NSDUs divided by the time between the first and last N-DATA indications in the sequence.

Successful transfer of the octets in a transmitted NSDU is defined to occur when the octets are delivered to the intended receiving NS user without error, in the proper sequence, prior to release of the NC by the receiving NS user.

Throughput is specified separately for each direction of transfer. Each Throughput specification will specify both the desired "target" value and the minimum acceptable value (i.e., the "lowest quality acceptable") for the NC. (See also 12.2.7.)

10.2.4 Transit delay

Transit delay is the elapsed time between an N-DATA request and the corresponding N-DATA indication. Elapsed time values are calculated only on NSDUs that are successfully transferred.

Successful transfer of an NSDU is defined to occur when the NSDU is transferred from the sending NS user to the intended receiving NS user without error, in the proper sequence, prior to release of the NC by the receiving NS user.

Specification of transit delay will define a pair of values: the desired "target" value and the maximum acceptable (i.e. the "lowest quality acceptable") value. (See also 12.2.7.) The specified values will be averages and will be based on an NSDU size of 128 octets.

The pair of transit delay values specified for an NC applies to both directions of transfer. That is, the transit delay in each direction is expected to be no worse than that specified.

The transit delay for an individual NSDU may be increased if the receiving NS user exercises flow control. Such occurrences are excluded in calculating both average and maximum transit delay values.

10.2.5 Residual error rate

Residual error rate is the ratio of total incorrect, lost, and duplicate NSDUs to total NSDUs transferred across the NS boundary during a measurement period. The relationship among these quantities is defined, for a particular NS user pair, as shown in figure 3.

10.2.6 Transfer failure probability

Transfer failure probability is the ratio of total transfer failures to total transfer samples observed during a performance measurement.

A transfer sample is a discrete observation of NS provider performance in transferring NSDUs between a specified sending and receiving NS user. A transfer sample begins on input of a selected NSDU at the sending NS user boundary, and continues until the outcome of a given number of NSDU transfer requests has been determined. A transfer sample will normally correspond to the duration of an individual NC.

A transfer failure is a transfer sample in which the observed performance is worse than a specified minimum acceptable level. Transfer failures are identified by comparing the measured values for the supported performance parameters with specified transfer failure thresholds. The three supported performance parameters are throughput, transit delay, and residual error rate.

In systems where Network Service QOS is reliably monitored by the NS provider, transfer failure probability can be estimated by the probability of an NS provider invoked N-DISCONNECT during a transfer sample.




Figure 3 — Components of residual error rate

10.2.7 NC resilience

NC resilience parameters specify the probability of

- a) an NS provider invoked NC release (i.e. issuance of an N-DISCONNECT indication with no prior N-DISCONNECT request); and
- b) an NS provider invoked reset (i.e., issuance of an N-RESET indication with no prior N-RESET request);

during a specified time interval on an established NC.

10.2.8 NC release delay

NC release delay is the maximum acceptable delay between an NS user invoked N-DISCONNECT request and the successful release of the NC at the peer NS user. NC release delay is normally specified independently for each NS user. NC release delay does not apply in cases where NC release is invoked by the NS provider.

Issuance of an N-DISCONNECT request by either NS user starts the counting of NC release delay for the other NS user. Successful NC release is signalled to the NS user not initiating the N-DISCONNECT request by an N-DISCONNECT indication.

10.2.9 NC release failure probability

NC release failure probability is the ratio of total NC release requests resulting in release failure to total NC release requests included in a measurement sample. NC release failure probability is normally specified independently for each NS user.

A release failure is defined to occur, for a particular NS user, if that user does not receive an N-DISCONNECT indication

within the specified maximum NC release delay of the NS user issuing the N-DISCONNECT request (given that the former NS user has not issued an N-DISCONNECT request).

10.2.10 NC protection

NC protection is the extent to which an NS provider attempts to prevent unauthorized masquerading or monitoring or manipulation of NS-user-data. NC protection for an NC is specified by selecting any combination of the following features:

- a) confidentiality of an entire NSDU sequence on the NC;
- b) detection of modification, deletion, replay, or insertion of data within the NSDU sequence on an NC;
- c) peer entity authentication. The NS user may request that the NS provider should confirm the identity of the remote NSAP such that there is protection against masquerading by T-entities;
- d) authentication of the origin of an NSDU such that there is protection against the unauthorized insertion or replay of the NSDU.

10.2.11 NC priority

NC priority specifies independently the relative importance of an NC with respect to the following:

- a) priority to gain an NC;
- b) priority to keep an NC;
- c) priority of data on the NC.

NC priority QOS-parameters a) and b) together define the order in which NCs are to be broken to recover resources if

necessary. The NS provider is required to accept new requests for NCs with a high priority type a) if it can, even if NCs with a lower priority type b) have to be released to do so.

NC priority QOS-parameter c) defines the order in which NCs are to have their QOS degraded. The NCs with a high priority type c) are to have their requests serviced within the required QOS first and remaining resources are then used to attempt to satisfy requests on lower priority NCs.

NOTE — The use or abuse of the NC priority QOS-parameters can be controlled by one or more of the following :

- user discipline within a closed group of NS users;
- differential tariffs;

— management facilities within the Network Layer such that requests for NC priority are policed and regulated.

10.2.12 Maximum acceptable cost

The maximum acceptable cost QOS-parameter specifies the maximum acceptable cost for an NC. The cost may be specified in absolute or relative cost units. The cost of an NC is composed of communications and end-system resource costs.

NOTE — The possible actions of the NS provider in the event that the maximum acceptable cost for an NC is exceeded are not specified in this International Standard.

Section two : Definition of the connection-mode primitives

11 Sequence of primitives

This clause defines the constraints on the sequences in which the primitives defined in clauses 12 to 14 may occur. The constraints determine the order in which primitives occur, but do not fully specify when they may occur. Other constraints, such as flow control of data, will affect the ability of an NS user or an NS provider to issue a primitive at any particular time.

Table 5 is a summary of the NS primitives and their parameters.

11.1 Relation of primitives at the two NC endpoints

A primitive issued at one NC endpoint will, in general, have consequences at the other NC endpoint. The relations of primitives of each type to primitives at the other NC endpoint are defined in the appropriate clauses 12 to 14; all these relations are summarized in the diagrams in figure 4.

However, an N-DISCONNECT request or indication primitive may terminate any of the other sequences before completion.

Table 5 — Summary of Network Service primitives and parameters

Phase	Service	Primitive	Parameters
NC establishment	NC establishment	N-CONNECT request	(Called address, calling address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT indication	(Called address, calling address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT response	(Responding address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
		N-CONNECT confirm	(Responding address, receipt confirmation selection, expedited data selection, QOS-parameter set, NS-user-data)
Data transfer	Data transfer	N-DATA request	(NS-user-data, confirmation request)
		N-DATA indication	(NS-user-data, confirmation request)
	Receipt confirmation (see the note)	N-DATA-ACKNOWLEDGE request	—
		N-DATA-ACKNOWLEDGE indication	—
	Expedited data transfer (see the note)	N-EXPEDITED-DATA request	(NS-user-data)
		N-EXPEDITED-DATA indication	(NS-user-data)
	Reset	N-RESET request	(Reason)
		N-RESET indication	(Originator, reason)
		N-RESET response	—
		N-RESET confirm	—
NC release	NC release	N-DISCONNECT request	(Reason, NS-user-data, responding address)
		N-DISCONNECT indication	(Originator, reason, NS-user-data, responding address)

NOTE — An NS provider-option service; it may not be provided in every Network Service.




Figure 4 — Summary of Network Service primitive time sequence diagrams

An N-RESET request or indication may terminate a data transfer, expedited data transfer, or receipt confirmation sequence before completion.

11.2 Sequence of primitives at one NC endpoint

The possible overall sequences of primitives at an NC endpoint are defined in the state transition diagram, figure 5. In the diagram

a) a primitive which is not shown as resulting in a transition (from one state to the same state, or from one state to a different state) is not permitted in that state (however, see 11.1 concerning the effect of N-DISCONNECT and N-RESET primitives);

b) N-DISCONNECT stands for either the request or the indication form of the primitive in all cases;

c) the labelling of the states NS user invoked reset pending (state 5) and NS provider invoked reset pending (state 6) indicates the party which started the local interaction, and does not necessarily reflect the value of the originator parameter in the associated N-RESET primitive;

d) the Idle (state 1) reflects the absence of an NC. It is the initial and final state of any sequence, and once it has been re-entered, the NC is released;

e) the use of a state transition diagram to describe the allowable sequences of service primitives does not impose any requirements or constraints on the internal organization of any implementations of the Network Service.




Figure 5 — State transition diagram for sequences of primitives at an NC endpoint

12 Network Connection establishment phase

12.1 Function

The NC establishment service primitives can be used to establish an NC, provided the NS users exist and are known to the NS provider.

Simultaneous N-CONNECT requests at the two NSAPs are handled independently by the NS provider; they may result in two, one or zero NCs.

12.2 Types of primitives and parameters

Table 6 indicates the types of primitives and the parameters needed for NC establishment.

12.2.1 Addresses

The parameters which take addresses as values (12.2.2 to 12.2.4) all refer to NSAP addresses. The NSAP address parameters will accommodate variable length addresses up to a defined maximum of 40 decimal digits (when expressed in decimal digit syntax).

The values of these addresses as supplied by the NS user are not necessarily checked or authenticated by the NS provider. An NS user receiving these addresses in N-CONNECT indication or confirm primitives can only rely on their validity if the NS user has knowledge that the NS provider guarantees address correctness.

NOTE — Mechanisms operating within the NS provider, such as call redirection or resolution of generic addresses, may result in address parameters in corresponding primitives not being identical in the following cases:

- the responding address parameter on the N-CONNECT response may not necessarily be the same as the called address parameter on the N-CONNECT indication;
- the responding address parameter on the N-CONNECT confirm may not necessarily be the same as the called address parameter on the N-CONNECT request.

12.2.2 Called address parameter

The called address parameter conveys an address identifying the NSAP to which the NC is to be established. Where explicitly supplied, the addresses in corresponding N-CONNECT request and indication primitives are identical.

12.2.3 Calling address parameter

The calling address parameter conveys the address of the NSAP from which the NC has been requested. Where explicitly supplied, the addresses in corresponding N-CONNECT request and indication primitives are identical.

12.2.4 Responding address parameter

The responding address parameter conveys the address of the NSAP to which the NC has been established. Where explicitly supplied, the addresses in corresponding N-CONNECT response and confirm primitives are identical. This parameter always conveys a specific NSAP address and not a generic NSAP address.

12.2.5 Receipt confirmation selection parameter

The receipt confirmation selection parameter indicates the use/availability of the receipt confirmation service on the NC. If the receipt confirmation service is not provided in the Network Service, then it cannot be used on the NC (see clause 8). The value of this parameter is either "use of receipt confirmation" or "no use of receipt confirmation". The values on the various primitives are related such that

- on the N-CONNECT request, either of the defined values may occur;
- on the N-CONNECT indication, the value is either equal to the value on the request primitive, or is "no use of receipt confirmation";
- on the N-CONNECT response, the value is either equal to the value on the indication primitive or is "no use of receipt confirmation";
- on the N-CONNECT confirm, the value is equal to the value on the response primitive.

Table 6 — NC establishment primitives and parameters

Parameter \ Primitive	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
Called address	X	X (=) (see the note)		
Calling address	X (see the note)	X (=)		
Responding address			X (see the note)	X (=)
Receipt confirmation selection	X	X	X	X (=)
Expedited data selection	X	X	X	X (=)
QOS-parameter set	X	X	X	X (=)
NS-user-data	X	X (=)	X	X (=)

NOTE — This parameter may be implicitly associated with the NSAP at which the primitive is issued.

ISO 8348 : 1987 (E)

Since receipt confirmation may not be provided in the Network Service and since, when it is available, both NS users and the NS provider must agree to its use, there are four possible cases of negotiation of receipt confirmation on an NC:

- a) the calling NS user does not request it — it is not used;
- b) the calling NS user requests it but the NS provider does not provide it — it is not used;
- c) the calling NS user requests it and the NS provider agrees to provide it, but the called NS user does not agree to its use — it is not used;
- d) the calling NS user requests it, the NS provider agrees to provide it, and the called NS user agrees to its use — it can be used.

12.2.6 Expedited data selection parameter

The expedited data selection parameter indicates the use/availability of the expedited data transfer service on the NC. If the expedited data transfer service is not available from the NS provider (see clause 8), then it cannot be used on the NC. The value of this parameter is either "use of expedited data" or "no use of expedited data". The values on the various primitives are related such that

- a) on the N-CONNECT request, either of the defined values may occur;
- b) on the N-CONNECT indication, the value is either equal to the value on the request primitive, or is "no use of expedited data";
- c) on the N-CONNECT response, the value is either equal to the value on the indication primitive or is "no use of expedited data";
- d) on the N-CONNECT confirm, the value is equal to the value on the response primitive.

12.2.7 QOS-parameter set

For each QOS-parameter which is conveyed during NC establishment, a set of "subparameters" is defined from among the following possibilities:

- a) a "target" value which is the QOS value desired by the calling NS user;
- b) the "lowest quality acceptable" value which is the lowest QOS value agreeable to the calling NS user;
- c) an "available" value which is the QOS value the NS provider is willing to provide; and
- d) a "selected" value which is the QOS value to which the called NS user agrees.

The set of values which can be specified for each subparameter is defined in every Network Service. Each set of values includes the value "unspecified". It may also include a value defined to be a "default value", which is mutually understood by the NS provider and the NS user between which it is conveyed.

NOTE — "Default" values are defined between a particular NS user and the NS provider. Different "defaults" may exist for different NS users and thus a value which is understood as a "default" at one end of an NC may not be the "default" value at the other end.

In those cases where both the subparameters "target" and "lowest quality acceptable" are specified by the calling NS user, they are boundary parameters defining a "range" of QOS values to which the calling NS user will agree. Similarly, where both the subparameters "available" and "lowest quality acceptable" are specified by the NS provider, they are boundary parameters defining a "range" of QOS values which the NS provider is willing to provide. These "ranges" are defined to include the values of both of the boundary subparameters, plus any values allowed for these subparameters which lie between the boundary subparameters. In the case where the "target" (or the "available") subparameter has a specified value but the "lowest quality acceptable" value is "unspecified", the range is defined to consist of the "target" value plus all other values which are allowed for these subparameters and which are lower (in QOS terms) than the target. If the value for both the "target" and "lowest quality acceptable" is "unspecified", then no range of values is defined.

NOTE — For other value assignments (for example "target" is "unspecified" but "lowest quality acceptable" has a specified value), the "range" is not defined since these assignments are not allowed in the negotiation procedures described in 12.2.7.1 and 12.2.7.2.

12.2.7.1 Throughput

Table 7 indicates the presence of the QOS-subparameters for the throughput QOS-parameters in the N-CONNECT primitives.

The negotiation and conveyance of each of the two throughput QOS-parameters are conducted as follows:

- a) In the N-CONNECT request primitive, the calling NS user specifies values for the "target" and "lowest quality acceptable" (i.e. lowest throughput) subparameters. Permitted value assignments are

Case 1: both the "target" and "lowest quality acceptable" are "unspecified";

Case 2: values other than "unspecified" are specified for both "target" and "lowest quality acceptable";

Case 3: a value other than "unspecified" is specified for the "target" and the "lowest quality acceptable" is "unspecified".

NOTE — The case where "target" is "unspecified" and the "lowest quality acceptable" has a value other than "unspecified" is not permitted; logically, this case can be represented by the permitted assignment where an identical value is specified for both the "target" and "lowest quality acceptable" (case 2).

- b) If the value assignments of the "target" and "lowest quality acceptable" subparameters are as defined in case 1, then the NS provider determines the highest QOS throughput value which is to be offered on the NC. This value (which may be the "default" value understood by the NS provider and the called NS user) is specified as the "available" subparameter in the N-CONNECT indication while the "lowest quality acceptable" subparameter value is "unspecified". If the requested QOS value assignments are as defined in case 2 or case 3, then, if the NS provider does not agree to provide a QOS in the requested range, the NC establishment attempt is rejected as described in 13.5. If the NS provider does agree to provide a QOS in the requested range, then in the N-CONNECT indication, the "available"

subparameter specifies the highest QOS value within the range which the NS provider is willing to provide and the "lowest quality acceptable" subparameter value is identical to that of the "lowest quality acceptable" subparameter in the N-CONNECT request.

c) If the called NS user does not agree to a QOS in the range between the "available" and the "lowest quality acceptable" subparameters of the N-CONNECT indication then the NS user rejects the NC establishment attempt as described in 13.4.

d) If the called NS user does agree to a QOS in the specified range, then the NS user specifies the agreed to value in the "selected" parameter of the N-CONNECT response.

e) In the N-CONNECT confirm, the "selected" subparameter has a value identical to that of "selected" in the N-CONNECT indication.

A summary of the negotiation procedures for the throughput QOS-subparameters is contained in table 8.

Table 7 – Negotiated QOS-subparameters for throughput QOS-parameters

Parameter	Primitive	N-CONNECT request	N-CONNECT indication	N-CONNECT response	N-CONNECT confirm
Throughput 1 "target" (calling to called)		X			
Throughput 1 "lowest quality acceptable" (calling to called)		X	X(=)		
Throughput 2 "target" (called to calling)		X			
Throughput 2 "lowest quality acceptable" (called to calling)		X	X(=)		
Throughput 1 "available" (calling to called)			X		
Throughput 2 "available" (called to calling)			X		
Throughput 1 "selected" (calling to called)				X	X(=)
Throughput 2 "selected" (called to calling)				X	X(=)

Table 8 – Negotiation of throughput QOS-subparameters

	Calling NS user specifies in N-CONNECT request		NS provider specifies in N-CONNECT indication		Called NS user specifies in N-CONNECT response	NS provider specifies in N-CONNECT confirm	Notes
	"Target"	"Lowest quality acceptable"	"Available"	"Lowest quality acceptable"	"Selected"	"Selected"	
Case 1	"Unspecified"	"Unspecified"	Z	"Unspecified"	A	A	Z may be a "default" value; $Z > A > 0$
Case 2	X	Y	Z	Y	A	A	X and/or Y may be defined to be the "default" value at the calling NS user end, called NS user end, or both; $X > Z > Y$; $Z > A > Y$
Case 3	X	"Unspecified"	Z	"Unspecified"	A	A	X may be a "default" value; $X > Z > 0$ $Z > A > 0$