
British Standard for
**Open Systems Interconnection:
Basic connection-oriented session service
definition**

[ISO title: Information processing systems — Open Systems Interconnection —
Basic connection oriented session service definition]

Interconnexion de Systèmes Ouverts: Service de session en mode connexion

Kommunikation Offener Systeme; Definition der verbindungsorientierten Basisdienste
der Kommunikationssteuerungsschicht

Contents

	Page		Page
National foreword	1	12 Typed data primitives and parameters	21
Committees responsible	Back cover	13 Capability data exchange primitives and parameters	21
Service Definition		14 Give tokens primitives and parameters	22
0 Introduction	2	15 Please tokens primitives and parameters	22
1 Scope and field of application	2	16 Minor synchronization point primitives and parameters	23
2 References	2	17 Major synchronization point primitives and parameters	24
Section one. General		18 Resynchronize primitives and parameters	25
3 Definitions	3	19 P-exception reporting primitives and parameters	26
4 Symbols and abbreviations	4	20 U-exception reporting primitives and parameters	27
5 Conventions	4	21 Activity start primitives and parameters	28
6 Model of the session service	4	22 Activity resume primitives and parameters	28
7 Overview of the session service	4	23 Activity interrupt primitives and parameters	29
8 Phases and services of the session service	6	24 Activity discard primitives and parameters	30
9 Functional units and subsets	7	25 Activity end primitives and parameters	30
10 Quality of session service	9	26 Orderly release primitives and parameters	31
Section two. Session service primitives		27 U-abort primitives and parameters	32
11 Introduction to session service primitives	14	28 P-abort primitives and parameters	32
12 Session connection establishment phase	18	29 Indications resulting from collision resolution	33
13 Data transfer phase	20	30 Events generated by the SS-provider	37
14 Session connection release phase	31	31 States	37
15 Sequences of primitives	32	32 Events generated by the SS-user	38
16 Collision	32	33 Operations on variables	39
Annexes		34 Specific actions	40
A State tables	34	35 Predicates	41
B Corrections incorporated in ISO 8326 : 1987	47	36 Connection establishment state table	41
Tables		37 Data transfer state table	42
1 Functional units using tokens	7	38 Synchronization state table	42
2 Services associated with each functional unit	8	39 Resynchronization state table	43
3 Classification of performance QOS parameters	10	40 Activity interrupt and discard state table	44
4 Parameters specifying other session service features	10	41 Activity start, resume and capability data state table	44
5 Session connection establishment phase primitives	14	42 Token management and exceptions state table	45
6 Data transfer phase primitives	15	43 Connection release state table	46
7 Session connection release phase primitives	16	Figures	
8 Token restrictions on service primitives	16	1 Relationship of this International Standard to other OSI standards	2
9 Session connection establishment primitives and parameters	19	2 Example of a structured dialogue unit	5
10 Normal data transfer primitives and parameters	20	3 Example of a structured activity	5
11 Expedited data transfer primitives and parameters	20	4 Components of residual error rate	13

National foreword

This British Standard has been prepared under the direction of the Information Systems Technology Standards Committee. It is identical with ISO 8326 : 1987 'Information processing systems — Open Systems Interconnection — Basic connection oriented session service definition', published by the International Organization for Standardization (ISO). It is a revision of DD 111 : 1985, which is withdrawn.

The UK has actively participated through BSI Technical Committee IST/21 in the development of the international standard.

The text for the international standard has been developed in conjunction with The International Telegraph and Telephone Consultative Committee (CCITT), and will be technically identical to CCITT Recommendation X.215 : 1988, which will be available from BSI Sales Department, Linford Wood, Milton Keynes MK14 6LE.

Draft addenda to the international standard covering session symmetric synchronization, unlimited user data and connectionless extensions are in preparation.

This British Standard is closely related to BS 6961, which specifies the protocol providing the service defined in this British Standard.

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

Terminology and conventions. The text of the international standard has been approved as suitable for publication as a British Standard without deviation. Some terminology and certain conventions are not identical with those used in British Standards; attention is drawn especially to the following.

Wherever the words 'International Standard' appear, referring to this standard, they should be read as 'British Standard'.

Cross-references

International standard	Corresponding British Standard
ISO 7498 : 1984	BS 6568 Reference model of Open Systems Interconnection Part 1 : 1988 Basic reference model (incorporating connectionless-mode transmission) (Identical)
ISO 8327 : 1987	BS 6961 : 1988 Open Systems Interconnection. Basic connection-oriented session protocol specification (Identical)
ISO/TR 8509 : 1987	DD 108 : 1988 Open Systems Interconnection: service conventions (Identical)

The Technical Committee has reviewed the provisions of ISO/DIS 7498-3 (referred to as ISO 7498/Add. 3) and ISO 8072, to which reference is made in the text, and has decided that they are acceptable for use in conjunction with this standard. A British Standard identical with ISO 8072 is in preparation.

NOTE. *Typographical error.* In clause 2 the title of ISO 7498/Add 3 should be 'Naming including addressing'.

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

British Standard for

Open systems Interconnection: Basic connection-oriented session service definition

0 Introduction

This International Standard is one of a set of International Standards produced to facilitate the interconnection of computer systems.

This International Standard is related to other International Standards in the set as defined by the Reference Model for Open Systems Interconnection (ISO 7498). The Reference Model subdivides the area of standardization for interconnection into a series of layers of specification, each of manageable size.

The purpose of this International Standard is to define the service provided to the Presentation Layer at the boundary between the Session and Presentation Layers of the Reference Model. The session service is provided by the session protocol making use of the services available from the Transport Layer. This International Standard also defines the session service characteristics which the presentation protocol may exploit. The relationship between the International Standards for the session service, session protocol, transport service, and the presentation protocol is illustrated in figure 1.

It is recognized that, with respect to session Quality of Service, (described in clause 10), work is still in progress to provide an integrated treatment of QOS across all of the layers of the OSI Reference Model and to ensure that the individual treatments in each layer service satisfy overall QOS objectives in a consistent manner. As a consequence, an addendum may be added to this International Standard at a later time which reflects further QOS developments and integration.

1 Scope and field of application

This International Standard defines in an abstract way the externally visible service provided by the OSI Session Layer in terms of

- the primitive actions and events of the service;
- the parameter data associated with each primitive action and event;
- the relationship between, and the valid sequence of these actions and events.

The service defined in this International Standard is that which is provided by the OSI session protocol (in conjunction with the transport service) and which may be used by the OSI presentation protocol.

This International Standard does not specify individual implementations or products, nor does it constrain the implementation of entities and interfaces within a computer system. There is, therefore, no conformance to this International Standard.

2 References

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

ISO 7498/Add.3, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Addendum 3: Name including addressing*.¹⁾

ISO 8072, *Information processing systems — Open Systems Interconnection — Transport service definitions*.

ISO 8327, *Information processing systems — Open Systems Interconnection — Basic connection oriented session protocol specification*.

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




Figure 1 — Relationship of this International Standard to other OSI standards

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

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-session-service-data-unit;
- b) session-connection;
- c) Session Layer;
- d) session-service;
- e) session-service-access-point;
- f) session-service-data-unit;
- g) Transport Layer;
- h) duplex;
- i) half-duplex.

3.2 Service convention definitions

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

- a) service-user;
- b) service-provider;
- c) primitive;
- d) request;
- e) indication;
- f) response;
- g) confirm.

3.3 Session-service definitions

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

3.3.1 calling SS-user: An SS-user that initiates a session connection establishment request.

3.3.2 called SS-user: An SS-user with whom a calling SS-user wishes to establish a session connection.

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

3.3.3 sending SS-user: An SS-user that acts as a source of data during the data transfer phase of a session connection.

3.3.4 receiving SS-user: An SS-user that acts as a sink of data during the data phase of a session connection.

NOTE — An SS-user can be both a sending and a receiving SS-user simultaneously.

3.3.5 requestor; requesting SS-user: An SS-user that initiates a particular action.

3.3.6 acceptor; accepting SS-user: An SS-user that accepts a particular action.

3.3.7 token: An attribute of a session connection which is dynamically assigned to one SS-user at a time to permit certain services to be invoked.

3.3.8 conditional (parameter): A parameter whose presence in a request or response depends on conditions defined in the text of this International Standard; and whose presence in an indication or confirm is mandatory if that parameter was present in the preceding session service primitive, or absent if that parameter was absent in the preceding session service primitive.

3.3.9 proposed parameter: The value for a parameter proposed by an SS-user in an S-CONNECT request or an S-CONNECT response that it wishes to use on the session connection.

3.3.10 selected parameter: The value for a parameter that has been chosen for use on the session connection.

4 Symbols and abbreviations

4.1 Abbreviations

SS	: session-service
SSAP	: session-service-access-point
SSDU	: session-service-data-unit
NSSDU	: normal-data-session-service-data-unit
TSSDU	: typed-data-session-service-data-unit
XSSDU	: expedited-session-service-data-unit
QOS	: quality of service

4.2 Service variables

V(A)	See 11.4.1.1
V(M)	See 11.4.1.2
V(R)	See 11.4.1.3
Vsc	See 11.4.1.4

5 Conventions

This International Standard uses the descriptive conventions defined in ISO/TR 8509 except that, where indicated in this International Standard, parameter values associated with a service primitive may be passed in a direction opposite to the direction of the service primitive.

6 Model of the session service

This International Standard uses the abstract model for a layer service defined in ISO/TR 8509. The model defines the interactions between the SS-user and the SS-provider which take place at the two SSAPs. Information is passed between an SS-user and the SS-provider by service primitives, which may convey parameters.

7 Overview of the session service

7.1 General overview

The session service provides the means for organized and synchronized exchange of data between cooperating SS-users. It provides its users with means to

- a) establish a connection with another SS-user, exchange data with that user in a synchronized manner, and release the connection in an orderly manner;
- b) negotiate for the use of tokens to exchange data, synchronize and release the connection, and to arrange for data exchange to be half-duplex or duplex;
- c) establish synchronization points within the dialogue and, in the event of errors, resume the dialogue from an agreed synchronization point;
- d) interrupt a dialogue and resume it later at a prearranged point.

7.2 Token concept

A token is an attribute of a session connection which is dynamically assigned to one SS-user at a time to permit certain services to be invoked. It is the right to exclusive use of the service.

Four tokens are defined:

- a) the data token;
- b) the release token;
- c) the synchronize-minor token;
- d) the major/activity token.

A token is always in one of the following states:

- e) available, in which case it is always
 - 1) assigned to one SS-user, who then has the exclusive right to use the associated service (provided that no other restrictions apply); and
 - 2) not assigned to the other SS-user, who does not have the right to use the service but may acquire it later;

or

- f) not available to either SS-user, in which case neither SS-user has the exclusive use of the associated service. The service then becomes inherently available to both SS-users (data transfer and release), or otherwise unavailable to both SS-users (synchronization and activities).

Restrictions related to the availability and assignment of tokens are defined in 11.2.

7.3 Synchronization and dialogue unit concepts

SS-users may insert synchronization points into the data they are transmitting. Each synchronization point is identified by a serial number maintained by the SS-provider (see 11.4).

Any semantics which SS-users may give to their synchronization points are transparent to the SS-provider.

There are two types of synchronization points:

- a) minor synchronization points;
- b) major synchronization points.

Major synchronization points are used to structure the exchange of data into a series of dialogue units. The characteristic of a dialogue unit is that all communication within it is completely separated from all communication before and after it. A major synchronization point indicates the end of one dialogue unit and the beginning of the next. Each major synchronization point is confirmed explicitly.

Minor synchronization points are used to structure the exchange of data within a dialogue unit. Figure 2 illustrates how a dialogue unit is structured through the use of minor synchronization points. Each minor synchronization point may or may not be confirmed explicitly.

7.4 Activity concept

The activity concept allows SS-users to distinguish between different logical pieces of work called activities. Each activity consists of one or more dialogue units. Only one activity is allowed on a session connection at a time, but there may be several consecutive activities during a session connection. An activity may also span more than one session connection.

An activity can be interrupted and then resumed on the same or on a subsequent session connection. This can be considered as a form of resynchronization.

Figure 3 shows how an activity may be structured into dialogue units through the use of major synchronization points. In addition, the SS-users may transfer data outside an activity.

7.5 Resynchronization

Resynchronization may be initiated by either SS-user. It sets the session connection to a defined state, and therefore includes reassignment of tokens and setting the synchronization point serial number to a new value. Resynchronization purges all undelivered data.




Figure 2 — Example of a structured dialogue unit




Figure 3 — Example of a structured activity

Three options are provided:

- a) abandon option which is used to set the synchronization point serial number to an unused value;
- b) restart option which is used to set the synchronization point serial number to any used value which is greater than the synchronization point serial number which identifies the last acknowledged major synchronization point;
- c) set option which is used to set the synchronization point serial number to any value chosen by the SS-user.

7.6 Negotiation

Negotiation takes place between both SS-users during the session connection establishment phase according to the following rules.

7.6.1 Negotiation of functional units

The kernel functional unit (see clause 9) is always used. Each SS-user proposes the use or non-use of each of the other functional units. A functional unit is selected only if both SS-users propose use of the functional unit and it is supported by the SS-provider. Specific negotiation rules are given in 12.1.2.

7.6.2 Negotiation of initial token settings

When the calling SS-user proposes use of a functional unit that requires a token, it also proposes the initial token settings:

- a) calling SS-user side;
- b) called SS-user side;
- c) called SS-user choice.

If the use of the functional unit is selected, the token is set to

- d) the side proposed by the called SS-user, if "called SS-user choice" is proposed by the calling SS-user; or
- e) in all other cases, the side proposed by the calling SS-user.

7.6.3 Negotiation of initial synchronization point serial number

When a calling SS-user proposes any of the major synchronize, minor synchronize or resynchronize functional units, but does not propose the activity management functional unit, it also proposes a value for the initial synchronization point serial number.

The calling SS-user may also propose a value for the initial synchronization point serial number even if the activity management functional unit is proposed provided that any of the minor synchronize, major synchronize or resynchronize functional units are also proposed. If the called SS-user selects use of any of the minor synchronize, major synchronize or resynchronize functional units, but does not select use of the activity management functional unit, it returns a value for the initial synchronization point serial number which may or may not be the same as the value proposed by the calling SS-user. The value returned by the called SS-user is used as the initial synchronization point serial number for the session connection.

In all other combinations of functional units, no initial synchronization point serial number is proposed.

8 Phases and services of the session service

The session service comprises three phases. The purpose of each phase, and a short description of the associated services is given in this clause. The services and the primitives by which they are invoked are defined in clauses 12, 13 and 14.

8.1 Session connection establishment phase

The session connection establishment phase is concerned with establishing a connection between two SS-users. It has one service associated with it:

the session connection service (see 12.1) is used to set up a session connection and to negotiate tokens and parameters to be used for the connection.

8.2 Data transfer phase

The data transfer phase is concerned with the exchange of data between the two SS-users connected in the session connection establishment phase.

There are four services associated with data transfer:

- a) the normal data transfer service (see 13.1) allows the transfer of normal data SSDUs (NSSDUs) over a session connection. Its use is controlled by the data token if the half-duplex functional unit has been selected;
- b) the expedited data transfer service (see 13.2) allows the transfer of expedited SSDUs (XSSDUs) over a session connection free from the token and flow control constraints of the normal data transfer service, typed data transfer service and capability data exchange service;
- c) the typed data transfer service (see 13.3) is used to transfer typed data SSDUs (TSSDUs) independent of the availability and assignment of the data token;
- d) the capability data exchange service (see 13.4) is used to exchange a limited amount of confirmed SS-user data while not within an activity.

There are three services concerned with token management:

- e) the give tokens service (see 13.5) allows an SS-user to surrender one or more specific tokens to the other SS-user;
- f) the please tokens service (see 13.6) allows an SS-user to request the other SS-user to transfer one or more specific tokens to it;
- g) the give control service (see 13.7) allows an SS-user to surrender all available tokens to the other SS-user.

There are three services associated with synchronization and resynchronization:

- h) the minor synchronization point service (see 13.8) allows the SS-user to separate the flow of NSSDUs and TSSDUs transmitted before the service was invoked from the subsequent flow of NSSDUs and TSSDUs. Its use is controlled by the synchronize-minor token;

i) the major synchronization point service (see 13.9) allows the SS-user to confine the flow of sequentially transmitted NSSDUs, TSSDUs and XSSDUs in each direction within a dialogue unit. Its use is controlled by the major/activity token;

j) the resynchronize service (see 13.10) is used to set the session connection to a previous or to a new synchronization point and to reassign the available tokens. This service may cause loss of NSSDUs, TSSDUs and XSSDUs.

There are two services for reporting errors or unanticipated situations:

k) the provider-initiated exception reporting service (see 13.11) (P-exception reporting service) permits SS-users to be notified of exception conditions or SS-provider protocol errors. This service may cause loss of NSSDUs, TSSDUs and XSSDUs;

l) the user-initiated exception reporting service (see 13.12) (U-exception reporting service) is used by the SS-user to report an exception condition when the data token is available but not assigned to the SS-user. This service may cause loss of NSSDUs, TSSDUs and XSSDUs.

There are five services associated with activities:

m) the activity start service (see 13.13) is used to indicate that a new activity is entered. Its use is controlled by the major/activity token;

n) the activity resume service (see 13.14) is used to indicate that a previously interrupted activity is re-entered. Its use is controlled by the major/activity token;

o) the activity interrupt service (see 13.15) allows an activity to be abnormally terminated with the implication that the work so far achieved is not to be discarded and may be resumed later. Its use is controlled by the major/activity token. This service may cause loss of NSSDUs, TSSDUs and XSSDUs;

p) the activity discard service (see 13.16) allows an activity to be abnormally terminated with the implication that the work so far achieved is to be discarded, and not resumed. Its use is controlled by the major/activity token. This service may cause loss of NSSDUs, TSSDUs and XSSDUs;

q) the activity end service (see 13.17) is used to end an activity (and set a major synchronization point). Its use is controlled by the major/activity token;

Using the activity services may lead to a state where no activity is in progress on the session connection. When activity services are employed, but no activity is in progress, only the activity start, activity resume, token management, capability data, typed data, normal data, expedited data, abort and release services may be invoked by the SS-users.

8.3 Session connection release phase

The session connection release phase is concerned with releasing a previously established session connection. It has three services associated with it:

a) the orderly release service (see 14.1) provides a means of achieving the orderly release of a session connection;

b) the user-initiated abort service (see 14.2) (U-abort service) is used to initiate the release of a session connection in a way that will terminate any outstanding service request. This service may cause loss of NSSDUs, TSSDUs and XSSDUs;

c) the provider-initiated abort service (see 14.3) (P-abort service) is used by the SS-provider to indicate the release of the session connection for internal reasons. This service may cause loss of NSSDUs, TSSDUs and XSSDUs. Any outstanding service request is terminated.

9 Functional units and subsets

9.1 Functional units

Functional units are logical groupings of related services defined by this International Standard for the purpose of

a) negotiation of SS-user requirements during the session connection establishment phase;

b) reference by other International Standards.

Table 1 specifies the association of tokens and functional units. When a functional unit implies the availability of a token, services concerned with the management of that token are provided in order to be able to request and transfer the available tokens.

The services associated with each functional unit are specified in table 2.

Table 1 — Functional units using tokens

Functional unit	Token
Negotiated release	Release token
Half-duplex	Data token
Minor synchronize	Synchronize-minor token
Major synchronize	Major/activity token
Activity management	Major/activity token

Table 2 — Services associated with each functional unit

Functional unit	Service(s)	Reference
Kernel (non-negotiable)	Session connection	12.1
	Normal data transfer	13.1
	Orderly release	14.1
	U-Abort	14.2
	P-Abort	14.3
Negotiated release	Orderly release	14.1
	Give tokens	13.5
	Please tokens	13.6
Half-duplex	Give tokens	13.5
	Please tokens	13.6
Duplex	No additional service	
Expedited data	Expedited data transfer	13.2
Typed data	Typed data transfer	13.3
Capability data exchange	Capability data exchange	13.4
Minor synchronize	Minor synchronization point	13.8
	Give tokens	13.5
	Please tokens	13.6
Major synchronize	Major synchronization point	13.9
	Give tokens	13.5
	Please tokens	13.6
Resynchronize	Resynchronize	13.10
Exceptions	Provider exception reporting	13.11
	User exception reporting	13.12
Activity management	Activity start	13.13
	Activity resume	13.14
	Activity interrupt	13.15
	Activity discard	13.16
	Activity end	13.17
	Give tokens	13.5
	Please tokens	13.6
	Give control	13.7

9.1.1 Kernel functional unit

The kernel functional unit supports the basic session services required to establish a session connection, transfer normal data and release the session connection.

9.1.2 Negotiated release functional unit

The negotiated release functional unit supports the negotiated orderly release service. The release token is available when this functional unit has been selected.

9.1.3 Half-duplex functional unit

The half-duplex functional unit supports the half-duplex service. The data token is available when this functional unit has been selected. It is not possible to select both this functional unit and the duplex functional unit for use on the same session connection.

9.1.4 Duplex functional unit

The duplex functional unit supports the duplex service. It is not possible to select both this functional unit and the half-duplex functional unit for use on the same session connection.

9.1.5 Expedited data functional unit

The expedited data functional unit supports the session expedited data transfer service.

9.1.6 Typed data functional unit

The typed data functional unit supports the typed data transfer service.

9.1.7 Capability data exchange functional unit

The capability data exchange functional unit supports the capability data exchange service. This functional unit can only be selected when the activity management functional unit has been selected.

9.1.8 Minor synchronize functional unit

The minor synchronize functional unit supports the minor synchronization point service. The synchronize-minor token is available when this functional unit has been selected.

9.1.9 Major synchronize functional unit

The major synchronize functional unit supports the major synchronization point service. The major/activity token is available when this functional unit has been selected.

9.1.10 Resynchronize functional unit

The resynchronize functional unit supports the resynchronize service.

9.1.11 Exceptions functional unit

The exceptions functional unit supports the user exception and provider exception reporting services.

This functional unit can only be selected when the half-duplex functional unit has been selected.

9.1.12 Activity management functional unit

The activity management functional unit supports the activity management services and the give control service. The major/activity token is available when this functional unit has been selected.

9.2 Subsets

A subset is defined as a combination of the kernel functional unit together with any other set of functional units provided that

- a) if the capability data functional unit is included in the subset, then the activity management functional unit is also included in the subset; and
- b) if the exceptions functional unit is included in the subset, then the half-duplex functional unit is also included in the subset.

NOTE — This International Standard contains no requirements for the registration of subsets. Users of this International Standard may define subsets to meet their session service needs. Other International Standards may identify subsets that comply with the above definition.

10 Quality of session service

The term "Quality of Service" (QOS) refers to certain characteristics of a session connection as observed between the session connection endpoints. QOS describes aspects of a session connection which are attributable solely to the SS-provider; such aspects are independent of SS-user behaviour (which is beyond the control of the SS-provider). SS-user behaviour does not impact the QOS provided.

Once a session connection is established, the SS-users at the two ends have the same knowledge and understanding of what the QOS over the session connection is.

10.1 Determination of QOS

QOS is described in terms of QOS parameters.

The definition of the QOS parameters associated with the session service is given in 10.3. These definitions provide both SS-users and the SS-provider with a common understanding of the QOS characteristics.

Two types of session service QOS parameters are identified:

- a) those which are negotiated during the session connection establishment phase:
 - 1) session connection protection (see 10.3.9);
 - 2) session connection priority (see 10.3.10);
 - 3) residual error rate (see 10.3.5);
 - 4) throughput, for each direction of transfer (see 10.3.3);
 - 5) transit delay, for each direction of transfer (see 10.3.4);
 - 6) optimized dialogue transfer (see 10.3.13); and
 - 7) extended control (see 10.3.12);
- b) those which are not negotiated during the session connection establishment phase but whose values are selected and/or known by other methods (for example, *a priori* knowledge and agreement, or by means of management functions) not defined in this International Standard:
 - 1) session connection establishment delay (see 10.3.1);
 - 2) session connection establishment failure probability (see 10.3.2);
 - 3) transfer failure probability (see 10.3.6);
 - 4) session connection release delay (see 10.3.7);
 - 5) session connection release failure probability (see 10.3.8);
 - 6) session connection resilience (see 10.3.11).

The negotiation procedures for parameters listed in a) are defined in 10.2. Once the session connection is established, the selected QOS parameters are not re-negotiated during the lifetime of the session connection. The SS-user should be aware that changes in QOS during a session connection are not signalled in the session service.

10.2 Session connection QOS negotiation procedures

QOS negotiation is described in terms of parameters which can be conveyed by the S-CONNECT primitives during the session connection establishment phase (see clause 12). For the parameters which are negotiated during the session connection establishment phase [see 10.1a)], the parameter values and negotiation rules are defined as follows:

- a) In the S-CONNECT request primitive, the calling SS-user can specify:
 - 1) for session connection protection, session connection priority, extended control, and optimized dialogue transfer, a single parameter value which is the "desired" QOS; for extended control and optimized dialogue transfer, one of the two values "feature desired" or "feature not desired" is conveyed;

NOTE — If the calling SS-user proposes use of the expedited data functional unit, the extended control parameter has the value "feature desired".

- 2) for residual error rate, and for each direction of throughput and transit delay, two parameter values which are the "desired" QOS and the "lowest acceptable" QOS to which the calling SS-user will agree.

b) In the S-CONNECT indication primitive, for each of the negotiated parameters, an "available" value is conveyed which is specified as follows:

1) for session connection protection, if the SS-provider agrees to provide a QOS value equivalent to the "desired" value specified in the S-CONNECT request, then the SS-provider specifies that value as "available"; if the SS-provider does not agree to provide the "desired" QOS requested, the SS-provider refuses to establish the session connection by issuing the S-CONNECT (reject) confirm primitive to the calling SS-user;

2) for session connection priority, the SS-provider specifies the QOS value it is willing to provide (a value which is equal to or better than the "desired" value specified in the S-CONNECT request) as "available";

3) for the residual error rate and each direction of throughput and transit delay, if the SS-provider agrees to provide a value of QOS which is equal to or better than the "lowest acceptable" QOS value specified in the S-CONNECT request, then the SS-provider specifies the value as "available"; if the SS-provider does not agree to provide this QOS, then the SS-provider refuses to establish the session connection by issuing the S-CONNECT (reject) confirm primitive to the calling SS-user;

4) for extended control and optimized dialogue transfer, if the "desired" value in the S-CONNECT request primitive is "feature not desired" then "feature not desired" is specified as "available"; if the "desired" value is "feature desired" and the SS-provider agrees to provide the feature on the session connection, then "feature desired" is specified as "available"; otherwise if the SS-provider does not agree to provide the feature, "feature not desired" is specified as "available".

c) In the S-CONNECT response primitive, for each of the negotiated parameters, an "agreed" value is conveyed which is specified as follows:

1) for optimized dialogue transfer, if the "available" value in the S-CONNECT indication primitive is "feature not desired" and the called SS-user agrees not to have the feature provided on the session connection, then "feature not desired" is specified as "agreed"; otherwise the SS-user may reject establishment of the session connection; if the "available" value in the indication primitive is "feature desired" and the SS-user agrees to have the feature provided, then "feature desired" is specified as "agreed"; otherwise, if the SS-user does not agree to provision of the feature, the value "feature not desired" is specified as "agreed";

2) for each of the other parameters, if the called SS-user agrees to the QOS value specified as "available" in the S-CONNECT indication primitive, then the identical value is specified as "agreed"; if the SS-user does not agree to the "available" value, the SS-user may reject establishment of the session connection.

d) In the S-CONNECT confirm primitive, for each of the negotiated parameters, an "agreed" value is conveyed which is identical to the "agreed" value conveyed in the S-CONNECT response.

10.3 Definition of QOS parameters

QOS parameters can be classified as

- a) parameters which express session service performance parameters, as shown in table 3;
- b) parameters which express other session service characteristics, as shown in table 4.

These session service QOS parameters are defined in this sub-clause.

Table 3 — Classification of performance QOS parameters

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

Table 4 — Parameters specifying other session service features

Extended control Session connection protection Session connection priority Optimized dialogue transfer

10.3.1 Session connection establishment delay

Session connection establishment delay is the maximum acceptable delay between an S-CONNECT request and the corresponding S-CONNECT confirm primitive.

NOTE — This delay includes SS-user dependent components.

10.3.2 Session connection establishment failure probability

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

Session connection establishment failure is defined to occur when a requested session connection is not established within the specified maximum acceptable session connection establishment delay as a result of misconnection, session connection refusal, or excessive delay on the part of the SS-provider. Session connection establishment attempts which fail as a result of error, session connection refusal, or excessive delay on the part of an SS-user are excluded in calculating session connection establishment failure probability.

10.3.3 Throughput

Throughput is defined for each direction of transfer, in terms of a sequence of at least two SSDUs successfully transferred by an S-DATA request/S-DATA indication or S-TYPED-DATA request/S-TYPED-DATA indication sequence of primitives. Given such a sequence of n SSDUs, where $n \geq 2$, the throughput is defined to be the smaller of

- a) the number of SS-user data octets contained in the last $n - 1$ SSDUs divided by the time between the first and last S-DATA or S-TYPED-DATA request in the sequence; and
- b) the number of SS-user data octets contained in the last $n - 1$ SSDUs divided by the time between the first and last S-DATA or S-TYPED-DATA indications in the sequence.

Successful transfer of the octets in a transmitted SSDU is defined to occur when the bits are delivered to the intended receiving SS-user without error, in the proper sequence, prior to release of the session connection by the receiving SS-user.

Throughput is only meaningful for a sequence of complete SSDUs and each specification is based on a previously stated average SSDU size.

Throughput is specified separately for each direction of transfer on a session connection. In each direction, a specification of throughput will consist of a "maximum throughput" value and an "average throughput" value. The "maximum throughput" value represents the maximum rate at which the SS-provider can continuously accept and deliver SSDUs, in the absence of sending SS-user input delays or flow control applied by the receiving SS-user. Thus, the sequence of SSDUs in the calculation above are defined to be presented continuously at the maximum rate. The "average throughput" value represents the expected transfer rate on a session connection including the effects of expected user-attributable delays (for example non-continuous SSDU input, receiving SS-user flow control). Thus, the sequence of SSDUs in the calculation above are defined to be presented at a rate which includes components representing "average" user delays.

It is possible for either the input or the output of a sequence of SSDUs to be excessively delayed by the SS-users. Such occurrences are excluded in calculating "average throughput" values.

For each direction of transfer, and for each of the "maximum throughput" and "average throughput" specifications, the throughput QOS for a particular session connection is negotiated between the SS-users and the SS-provider (see 10.2).

Throughput on a session connection relates only to the transfer of normal data and typed data over the session connection. There is no specification of the throughput for data which is transferred in association with the issue of any other session service primitives (for example S-CONNECT, S-CAPABILITY-DATA, etc.).

10.3.4 Transit delay

Transit delay is the elapsed time between the completion of any session service request primitive and the corresponding session service indication primitive occurring during the data transfer phase of a session connection. Elapsed time values are calculated only on service primitive pairs which are successfully completed.

Successful completion of a service primitive pair is defined to occur when the issue of the request primitive by one SS-user results in the issue of the corresponding indication primitive to the peer user (including any SS-user data associated with the primitive) which is without error, and in a proper sequence with respect to other primitives, prior to release of the session connection by the receiving SS-user.

In duplex and half-duplex session connections, transit delay is specified independently for each direction of transfer. In general, each transit delay specification defines both the average value and the maximum value expected for a session connection. Each specification of transit delay assumes a previously stated average size for SS-user data included in the service primitive pair.

An attempt to measure the transit delay for an individual service primitive pair may be greatly influenced if the receiving SS-user exercises flow control. Such occurrences are excluded in calculating both average and maximum transit delay values.

10.3.5 Residual error rate

Residual error rate is the ratio of total incorrect, lost, and duplicate units of SS-user data to the total units of SS-user data transferred across the session service boundary in association with any SS-primitive issued in the data transfer phase of a session connection during a measurement period. The relationship between these quantities for a particular SS-user pair is defined in figure 4.

10.3.6 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 SS-provider performance in handling service requests made by the SS-user. A transfer sample begins with the initiation of session service requests during the data transfer phase and continues until the outcome of a given number of service requests have been determined. These service requests may include the transfer of SS-user data or other service requests (such as S-ACTIVITY-START request, S-TOKEN-PLEASE request, etc.) made by the SS-user. A transfer sample will normally correspond to the duration of an individual session connection.

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 applicable supported performance parameters with specified transfer failure thresholds. The three supported performance parameters which may apply are throughput, transit delay, and residual error rate.

In systems where session service QOS is reliably monitored by the SS-provider, transfer failure probability can be estimated by the probability of an S-P-ABORT or an S-P-EXCEPTION-REPORT during a transfer sample.

10.3.7 Session connection release delay

Session connection release delay is the maximum acceptable delay between an SS-user initiated S-U-ABORT request and the successful release of a particular session connection. Session connection release delay is normally specified independently for each SS-user.

Issue of an S-U-ABORT request by either SS-user marks the beginning of the session connection release delay for both users. Successful release for one SS-user is defined to occur when that SS-user is first able to initiate a new session connection. Successful release is signalled to the SS-user not initiating the S-U-ABORT request by an S-U-ABORT indication. The SS-user initiating the S-U-ABORT request will normally receive a similar signal of local significance.

10.3.8 Session connection release failure probability

Session connection release failure probability is the ratio of total SS-user initiated abort requests resulting in session connection release failure to total SS-user initiated abort requests included in a measurement sample. Session connection release failure probability is normally specified independently for each SS-user.

Session connection release failure is defined to occur, for a particular SS-user, if that SS-user is not successfully released (as defined in 10.3.7) within the specified maximum session connection release delay as a result of error or excessive delay on the part of the SS-provider. Session connection release attempts which fail as a result of error, release refusal, or excessive delay on the part of an SS-user are excluded in calculating session connection release failure probability.

10.3.9 Session connection protection

Session connection protection is the extent to which an SS-provider attempts to prevent unauthorized monitoring or manipulation of SS-user originated information. Session connection protection is specified qualitatively by selecting one of the following session connection protection options:

- a) no protection features;
- b) protection against passive monitoring;
- c) protection against modification, replay, addition or deletion;
- d) both b) and c).

10.3.10 Priority

The specification of priority is concerned with the relationship between session connections. This parameter specifies the relative importance of a session connection with respect to

- a) the order in which session connections are to have their QOS degraded, if necessary; and
- b) the order in which session connections are to be broken to recover resources, if necessary.

This parameter only has meaning in the context of some management entity or structure able to judge relative importance. The number of priority levels is limited.

10.3.11 Session connection resilience

Session connection resilience parameters specify the probability of

- a) an SS-provider initiated non-orderly release of a session connection (i.e. issue of an S-P-ABORT indication); and
- b) an SS-provider exception report (i.e. issue of an S-P-EXCEPTION-REPORT indication) during a specified time interval on an established session connection.

10.3.12 Extended control parameter

The extended control parameter allows the SS-users to make use of the resynchronize, abort, activity interrupt and activity discard services when normal flow is congested.

NOTE — When the expedited data functional unit has been selected the extended control QOS is always provided to the SS-users.

10.3.13 Optimized dialogue transfer

The optimized dialogue transfer QOS parameter permits the concatenated transfer of certain session service requests. How this concatenation of service requests is achieved is a local implementation matter.

NOTE — This QOS parameter invokes the SS-provider extended concatenation protocol option.




Figure 4 — Components of residual error rate

Section two : Session service primitives

11 Introduction to session service primitives

11.1 Summary of primitives

Each of the services constituting the session service is achieved by invoking a sequence of session service primitives. Tables 5, 6 and 7 summarize the primitives and their parameters occurring in each phase of the session service. The parameters are defined in clauses 12, 13 and 14.

11.2 Token restrictions on sending primitives

Table 8 defines the conditions under which those service primitives requiring tokens may be issued.

11.3 Sequencing of primitives

All SS-user requests and responses are delivered by the SS-provider in the order in which they are submitted by the SS-user, except for the following:

- a) S-EXPEDITED-DATA;
- b) S-RESYNCHRONIZE;
- c) S-ACTIVITY-INTERRUPT;
- d) S-ACTIVITY-DISCARD;
- e) S-U-ABORT,

which may be delivered earlier than previously submitted primitives, but not later than subsequently submitted primitives.

Table 5 — Session connection establishment phase primitives

Service	Primitives	Parameters
Session Connection	S-CONNECT request S-CONNECT indication S-CONNECT response S-CONNECT confirm	Session Connection Identifier, Calling/ Called/ Responding Session Addresses, Result, QOS, Session Requirements, Synchronization Point Serial Number, Initial Assignment of Tokens, SS-user data