INTERNATIONAL STANDARD

ISO/IEC 10589

First edition
1992-06-15

Information technology — Telecommunications and information exchange between systems — Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)

Technologies de l'information — Communication de données et échange d'informations entre systèmes — Protocole intra-domaine de routage d'un système intermédiaire à un système intermédiaire à utiliser conjointement avec le protocole fournissant le service de réseau en mode sans connexion (ISO 8473)

Reference number
ISO/IEC 10589:1992(E)

Contents

ISO/IEC Copyright Office • Case Postale 56 • CH-1211 Genève 20 • Switzerland
Printed in Switzerland

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work.

In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1. Draft International Standards adopted by the joint technical committee are circulated to national bodies for voting. Publication as an International Standard requires approval by at least 75 % of the national bodies casting a vote.

International Standard ISO/IEC 10589 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*.

Annexes A and E form an integral part of this International Standard. Annexes B, C, D and F are for information only.

Introduction

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

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

The protocol is designed to operate in close conjunction with ISO 9542 and ISO 8473. ISO 9542 is used to establish connectivity and reachability between End Systems and Intermediate systems on individual subnetworks. Data is carried using the protocol specified in ISO 8473. The related algorithms for route calculation and maintenance are also described.

The intra-domain IS–IS routeing protocol is intended to support large routeing domains consisting of combinations of many types of subnetworks. This includes point-to-point links, multipoint links, X.25 subnetworks, and broadcast subnetworks such as ISO 8802 LANs.

In order to support large routeing domains, provision is made for Intra-domain routeing to be organised hierarchically. A large domain may be administratively divided into *areas*. Each system resides in exactly one area. Routeing within an area is referred to as *Level 1 routeing*. Routeing between areas is referred to as *Level 2 routeing*. Level 2 Intermediate systems keep track of the paths to destination areas. Level 1 Intermediate systems keep track of the routeing within their own area. For an NPDU destined to another area, a Level 1 Intermediate system sends the NPDU to the nearest level 2 IS in its own area, regardless of what the destination area is. Then the NPDU travels via level 2 routeing to the destination area, where it again travels via level 1 routeing to the destination End system.

Information technology — Telecommunications and information exchange between systems — Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)

1 Scope

This International Standard specifies a protocol which is used by Network Layer entities operating the protocol specified in ISO 8473 in Intermediate Systems to maintain routeing information for the purpose of routeing within a single routeing domain. The protocol specified in this International Standard relies upon the provision of a connectionless-mode underlying service.[1)]

This International Standard specifies:

a) procedures for the transmission of configuration and routeing information between network entities residing in Intermediate Systems within a single routeing domain;

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

c) procedures for the correct interpretation of protocol control information; and

d) the functional requirements for implementations claiming conformance to this International Standard.

The procedures are defined in terms of

e) the interactions between Intermediate system Network entities through the exchange of protocol data units;

f) the interactions between a Network entity and an underlying service provider through the exchange of subnetwork service primitives; and

g) the constraints on route determination which must be observed by each Intermediate system when each has a routeing information base which is consistent with the others.

2 Normative references

The following International Standards contain provisions which, through reference in this text, constitute provisions of this International Standard. At the time of publication, the editions indicated were valid. All International Standards are subject to revision, and parties to agreements based on this International Standard are encouraged to investigate the possibility of applying the most recent editions of the International Standards listed below. Members of IEC and ISO maintain registers of currently valid International Standards.

ISO 7498:1984, *Information processing systems — Open Systems Interconnection — Basic Reference Model.*

ISO 7498/Add.1:1987, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Addendum 1: Connectionless-mode Transmission.*

ISO 7498-3:1989, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Part 3: Naming and addressing.*

ISO 7498-4:1989, *Information processing systems — Open Systems Interconnection — Basic Reference Model — Part 4: Management framework.*

ISO/IEC 8208:1990, *Information technology — Data communications — X.25 packet Layer Protocol for Data Terminal Equipment.*

ISO 8348:1987, *Information processing systems — Data communications — Network service definition.*

ISO 8348/Add.1:1987, *Information processing systems — Data communications — Network Service Definition — Addendum 1: Connectionless-mode transmission.*

ISO 8348/Add.2:1988, *Information processing systems — Data communications — Network Service Definition — Addendum 2: Network layer addressing.*

ISO 8473:1988, *Information processing systems — Data communications — Protocol for providing the connectionless-mode network service.*

ISO/IEC 8473/Add.3:1989, *Information processing systems — Data Communications — Protocol for providing the connectionless-mode network service — Addendum 3: Provision of the underlying service assumed by ISO 8473 over subnetworks which provide the OSI data link service.*

ISO 8648:1990, *Information processing systems — Open Systems Interconnection — Internal organisation of the Network Layer.*

1) See ISO 8473 and its addendum 3 for the mechanisms necessary to realise this service on subnetworks based on ISO/IEC 8208, ISO 8802, and the OSI Data Link Service.

ISO/IEC 8802-1:[1)], *Information technology — Telecommunications and information exchange between systems — Local area networks — Part 1: General Introduction.*

ISO 8802-2:1989, *Information processing systems – Local area networks – Part 2: Logical link control.*

ISO/IEC 8802-3:1990, *Information processing systems — Local area networks — Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.*

ISO/IEC 8802-5:[1)], *Information technology — Local area networks — Part 5: Token ring access method and physical layer specifications.*

ISO/IEC 8802-6:[1)], *Information technology — Local area networks — Part 6: Distributed Queue Dual Bus (DQDB) access method and physical layer specifications.*

ISO/IEC 9314:1989, *Information processing systems — Fiber Distributed Data Interface (FDDI).*

ISO 9542:1988, *Information processing systems – Telecommunications and information exchange between systems — End system to Intermediate system Routeing exchange protocol for use in conjunction with the protocol for providing the connectionless -mode network service (ISO 8473).*

ISO/IEC TR 9575:1990, *Information technology — Telecommunications and information exchange between systems — OSI Routeing Framework.*

ISO/IEC TR 9577:1990, *Information technology — Telecommunications and information exchange between systems — Protocol identification in the network layer.*

ISO/IEC 10039:1991, *Information technology — Open Systems Interconnection — Local area networks — Medium Access Control (MAC) service definition.*

ISO/IEC 10165-1:[1)], *Information technology — Open Systems Interconnection — Structure of Management Information - Part 1: Management Information Model.*

ISO/IEC 10165-4:[1)], *Information technology — Open Systems Interconnection — Structure of management information — Part 4: Guidelines for the definition of managed objects.*

ISO/IEC 10733:[1)], *Information technology — Telecommunications and information exchange between systems — Elements of management information relating to OSI Network Layer standards.*

1) To be published

3 Definitions

3.1 Reference model definitions

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

a) Network Layer
b) Network Service access point
c) Network Service access point address
d) Network entity
e) Routeing
f) Network protocol
g) Network relay
h) Network protocol data unit

3.2 Network layer architecture definitions

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

a) Subnetwork
b) End system
c) Intermediate system
d) Subnetwork service
e) Subnetwork Access Protocol
f) Subnetwork Dependent Convergence Protocol
g) Subnetwork Independent Convergence Protocol

3.3 Network layer addressing definitions

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

a) Subnetwork address
b) Subnetwork point of attachment
c) Network Entity Title

3.4 Local area network definitions

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

a) Multi-destination address
b) Media access control
c) Broadcast medium

3.5 Routeing framework definitions

This International Standard makes use of the following terms defined in ISO/IEC TR 9575:

a) Administrative Domain
b) Routeing Domain
c) Hop
d) Black hole

3.6 Additional definitions

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

3.6.1 area: A routeing subdomain which maintains detailed routeing information about its own internal composition, and also maintains routeing information which allows it to reach other routeing subdomains. It corresponds to the Level 1 subdomain.

3.6.2 neighbour: An adjacent system reachable by traversal of a single subnetwork by a PDU.

3.6.3 adjacency: A portion of the local routeing information which pertains to the reachability of a single neighbour ES or IS over a single circuit.

Adjacencies are used as input to the Decision Process for forming paths through the routeing domain.

A separate adjacency is created for each neighbour on a circuit, and for each level of routeing (i.e. level 1 and level 2) on a broadcast circuit.

3.6.4 circuit: A subset of the local routeing information base pertinent to a single local SNPA. The system management view of a circuit is presented in a linkage managed object.

3.6.5 link: The communication path between two neighbours.

A link is "up" when communication is possible between the two SNPAs.

3.6.6 designated IS: The Intermediate system on a LAN which is designated to perform additional duties. In particular it generates Link State PDUs on behalf of the LAN, treating the LAN as a pseudonode.

3.6.7 pseudonode: Where a broadcast subnetwork has *n* connected Intermediate systems, the broadcast subnetwork itself is considered to be a pseudonode.

The pseudonode has links to each of the *n* Intermediate and End systems. Each of the ISs has a single link to the pseudonode (rather than *n–1* links to each of the other Intermediate systems). Link State PDUs are generated on behalf of the pseudonode by the Designated IS. This is depicted below in figure 1.

3.6.8 broadcast subnetwork: A subnetwork which supports an arbitrary number of End systems and Intermediate systems and additionally is capable of transmitting a single SNPDU to a subset of these systems in response to a single SN_UNITDATA request.

3.6.9 general topology subnetwork: A subnetwork which supports an arbitrary number of End systems and Intermediate systems, but does not support a convenient multi-destination connectionless transmission facility, as does a broadcast subnetwork.

3.6.10 routeing subdomain: a set of Intermediate systems and End systems located within the same Routeing domain.

3.6.11 level 2 subdomain: the set of all Level 2 Intermediate systems in a Routeing domain.

3.6.12 jitter: a small random variation introduced into the value of a timer to prevent multiple timer expirations in different systems from becoming synchronised.

4 Symbols and abbreviations

4.1 Data units

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

4.2 Protocol data units

ESH PDU	ISO 9542 End System Hello Protocol Data Unit
ISH PDU	ISO 9542 Intermediate System Hello Protocol Data Unit
RD PDU	ISO 9542 Redirect Protocol Data Unit
IIH PDU	Intermediate system to Intermediate system Hello Protocol Data Unit
LSP	Link State Protocol Data Unit
SNP	Sequence Numbers Protocol Data Unit
CSNP	Complete Sequence Numbers Protocol Data Unit
PSNP	Partial Sequence Numbers Protocol Data Unit

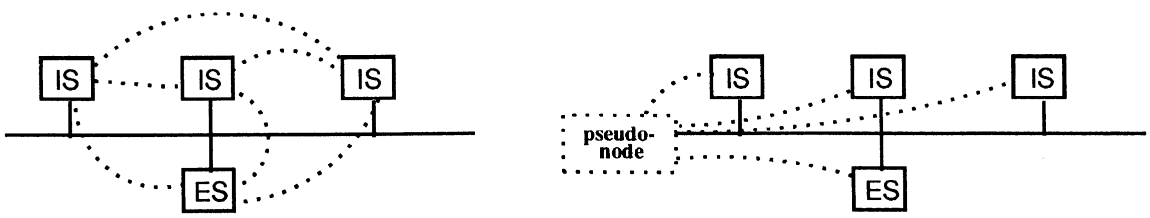

Figure 1 - Use of a pseudonode to collapse a LAN Topology

4.3 Addresses

AFI	Authority and Format Indicator
DSP	Domain Specific Part
IDI	Initial Domain Identifier
IDP	Initial Domain Part
NET	Network Entity Title
NPAI	Network Protocol Addressing Information
NSAP	Network Service Access Point
SNPA	Subnetwork Point of Attachment

4.4 Miscellaneous

DA	Dynamically Assigned
DED	Dynamically Established Data link
DTE	Data Terminal Equipment
ES	End System
IS	Intermediate System
HDLC	High Level Data Link Control
ISDN	Integrated Services Digital Network
FDDI	Fiber Distributed Data Interface
L1	Level 1
L2	Level 2
LAN	Local Area Network
MAC	Media Access Control
MAN	Metropolitan Area Network
NLPID	Network Layer Protocol Identifier
PSTN	Public Switched Telephone Network
OSIE	Open Systems Interconnection Environment
PCI	Protocol Control Information
QoS	Quality of Service
SN	Subnetwork
SNAcP	Subnetwork Access Protocol
SNDCP	Subnetwork Dependent Convergence Protocol
SNICP	Subnetwork Independent Convergence Protocol
SRM	Send Routeing Message
SSN	Send Sequence Numbers
SVC	Switched Virtual Circuit

5 Typographical conventions

This International Standard makes use of the following typographical conventions:

a) important terms and concepts appear in *italic* type when introduced for the first time;

b) protocol constants and management parameters appear in sansSerif type with multiple words run together. The first word is lower case, with the first character of subsequent words capitalised;

c) protocol field names appear in San Serif type with each word capitalised; and

d) values of constants, parameters, and protocol fields appear enclosed in "double quotes".

6 Overview of the protocol

6.1 System types

For the purposes of this International Standard, systems are classified according to the following types:

End Systems: These systems deliver NPDUs to other systems and receive NPDUs from other systems, but do not relay NPDUs. This International Standard does not specify any additional End system functions beyond those supplied by ISO 8473 and ISO 9542.

Level 1 Intermediate Systems: These systems deliver and receive NPDUs from other systems, and relay NPDUs from other source systems to other destina-

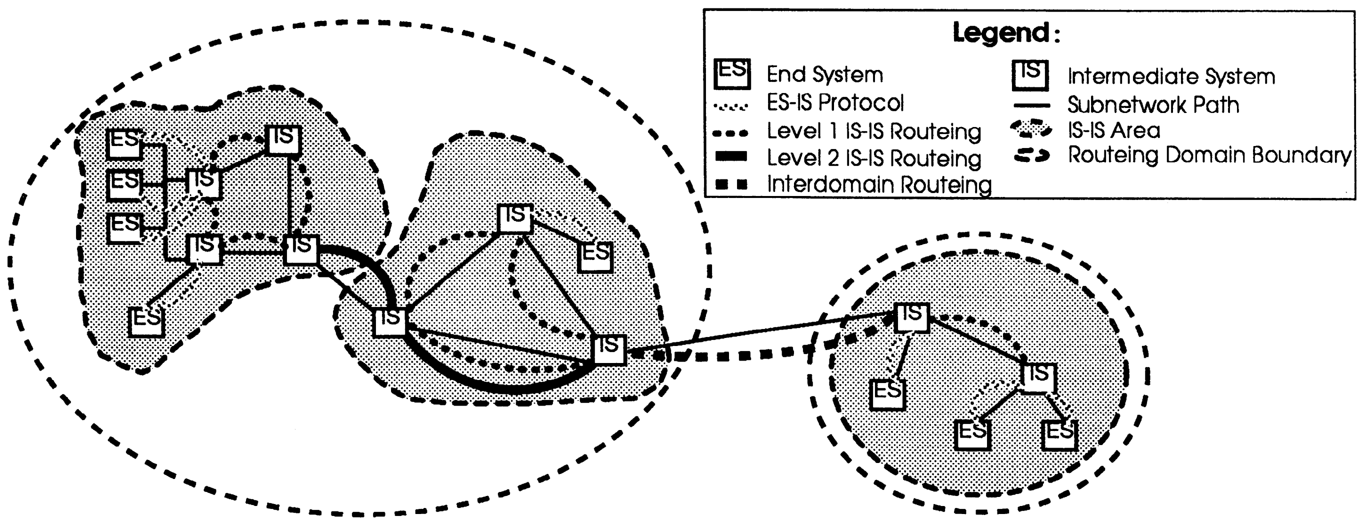

Figure 2 - Topologies and Systems supported by Intradomain Routeing

tion systems. They route directly to systems within their own area, and route towards a level 2 Intermediate system when the destination system is in a different area.

Level 2 Intermediate Systems: These systems act as Level 1 Intermediate systems in addition to acting as a system in the subdomain consisting of level 2 ISs. Systems in the level 2 subdomain route towards a destination area, or another routeing domain.

These systems and their topological relationship are illustrated in figure 2.

6.2 Subnetwork types

For the purposes of this International Standard, subnetworks are classified according to the following types:

a) *broadcast subnetworks*: These are multi-access subnetworks that support the capability of addressing a group of attached systems with a single NPDU, for instance ISO 8802-3 LANs.

b) *general topology subnetworks*: These are modelled as a set of point-to-point links each of which connects exactly two systems.

There are several generic types of general topology subnetworks:

1) *multipoint links*: These are links between more than two systems, where one system is a primary system, and the remaining systems are secondary (or slave) systems. The primary is capable of direct communication with any of the secondaries, but the secondaries cannot communicate directly among themselves.

2) *permanent point-to-point links*: These are links that stay connected at all times (unless broken, or turned off by system management), for instance leased lines or private links.

3) *dynamically established data links (DEDs):* These are links over connection oriented facilities, for instance X.25, X.21, ISDN, or PSTN networks.

Dynamically established data links can be used in one of two ways:

i) *static point-to-point (Static):* The call is established upon system management action and cleared only on system management action (or failure).

ii) *dynamically assigned (DA):* The call is established upon receipt of traffic, and brought down on timer expiration when idle. The address to which the call is to be established is determined dynamically from information in the arriving NPDU(s). No IS–IS routeing PDUs are exchanged between ISs on a DA circuit.

All subnetwork types are treated by the Subnetwork Independent functions as though they were connectionless subnetworks, using the Subnetwork Dependent Convergence functions of ISO 8473 where necessary to provide a connectionless subnetwork service. The Subnetwork Dependent functions do, however, operate differently on connectionless and connection-oriented subnetworks.

6.3 Topologies

A single organisation may wish to divide its *Administrative Domain* into a number of separate *Routeing Domains*. This has certain advantages, as described in ISO/IEC TR 9575. Furthermore, it is desirable for an intra-domain routeing protocol to aid in the operation of an inter-domain routeing protocol, where such a protocol exists for interconnecting multiple routeing domains.

In order to facilitate the construction of such multi-domain topologies, provision is made for the entering of inter-domain routeing information. This information is in the form of a set of *Reachable Address Prefixes* which may be entered either by System Management, or provided by an inter-domain routeing protocol at the ISs which have links crossing routeing domain boundaries. The prefix indicates that any NSAPs whose NSAP address matches the prefix may be reachable via the SNPA with which the prefix is associated. Where this SNPA is connected to a multi-destination subnetwork (e.g., dynamically assigned DED, broadcast), the prefix also has associated with it the required subnetwork addressing information, or an indication that it may be derived from the destination NSAP address (for example, an X.121 DTE address may sometimes be obtained from the IDI of the NSAP address).

The Address Prefixes are handled by the level 2 routeing algorithm in the same way as information about a level 1 area within the domain. NPDUs with a destination address matching any of the prefixes present on any Level 2 Intermediate System within the domain can therefore be relayed (using level 2 routeing) by that IS and delivered out of the domain. (It is assumed that the routeing functions of the other domain will then be able to deliver the NPDU to its destination.)

Where multiple routeing domains are interconnected using this International Standard, the model used is one in which the boundaries between routeing domains are on the subnetworks which connect the Intermediate systems. A boundary for a routeing domain is constructed by marking the linkage managed object associated with a circuit as being externalDomain rather than internal.

NOTE 1 This model also permits the construction of routeing domains whose scope is not limited by the hierarchical nature of network layer address assignment. For example, it is possible to construct a routeing domain, or even a single area, whose area addresses are taken from multiple addressing authorities.

6.4 Addresses

Within a routeing domain that conforms to this International Standard, the Network entity titles of Intermediate systems must meet the requirements stated in 7.1.4. It is the routeing domain administrative authority's responsibility to ensure that such is the case.

All systems shall be able to generate and forward NPDUs containing NSAP addresses in any of the formats specified by ISO 8348/Add.2. However, the routeing domain's administrative authority should ascertain that NSAP addresses of End systems meet the requirements set forth in 7.1.4 in order to take full advantage the routes derived by this protocol. Within such a domain it is still possible for some End systems to have addresses assigned which do not conform to the rules set forth in 7.1.4 provided that they meet the more general requirements of ISO 8348/Add.2, but these End systems may require additional configuration information to be entered into the Intermediate systems and they may obtain inferior routeing performance.

NOTE 2 The procedures whereby the routeing domain administrative authority obtains from an appropriate address authority Intermediate system NETs as required by this International Standard, and End system NSAP addresses as recommended by this International Standard are outside its scope.

6.5 Functional organisation

The intra-domain IS–IS routeing functions are divided into two groups

- Subnetwork Independent Functions
- Subnetwork Dependent Functions

6.5.1 Subnetwork independent functions

The Subnetwork Independent Functions supply full-duplex NPDU transmission between any pair of neighbour systems. They are independent of the specific subnetwork or data link service operating below them, except for recognising two generic types of subnetworks:

- **General Topology Subnetworks,** which include HDLC point-to-point, HDLC multipoint, and dynamically established data links (such as X.25, X.21, and PSTN links), and
- **Broadcast Subnetworks**, which include ISO 8802 LANs.

NOTE 3 This protocol is intended to operate on any broadcast subnetwork which meets the general requirements listed in 6.7. However, the remainder of this International Standard specifically addresses ISO 8802 LANs. Other LANs, such as FDDI, are believed to be adequately covered by the specification for ISO 8802 LANs. Other broadcast subnetworks, such as ISO 8802-6 MANs, may not be adequately covered at this time.

The following Subnetwork Independent Functions are identified:

- **Routeing.** The routeing function determines NPDU paths. A path is the sequence of connected systems and links between a source ES and a destination ES.

  The combined knowledge of all the Network Layer entities of all the Intermediate systems within a routeing domain is used to ascertain the existence of a path, and route the NPDU to its destination. The routeing component at an Intermediate system has the following specific functions:

  - It extracts and interprets the routeing PCI in an NPDU.
  - It performs NPDU forwarding based on the destination address.
  - It manages the characteristics of the path. If a system or link fails on a path, it finds an alternate route.
  - It interfaces with the subnetwork dependent functions to receive reports concerning an SNPA which has become unavailable, a system that has failed, or the subsequent recovery of an SNPA or system.
  - It informs the ISO 8473 error reporting function when the forwarding function cannot relay an NPDU, for instance when the destination is unreachable or when the NPDU would have needed to be segmented and the NPDU requested "no segmentation".
- **Congestion control.** Congestion control manages the resources used at each Intermediate system.

6.5.2 Subnetwork dependent functions

The subnetwork dependent functions mask the characteristics of the subnetwork or data link service from the subnetwork independent functions. These include:

- Operation of the Intermediate system functions of ISO 9542 on the particular subnetwork, in order to
  - determine neighbour Network entity title(s) and SNPA address(es);
  - determine the SNPA address(es) of operational Intermediate systems.
- Operation of the requisite Subnetwork Dependent Convergence Function as defined in ISO 8473 and its addendum 3, in order to perform
  - data link initialisation;
  - hop by hop fragmentation over subnetworks with small maximum SNSDU sizes; and
  - call establishment and clearing on dynamically established data links.

6.6 Design goals and non-goals

6.6.1 Goals

This International Standard supports the following design requirements. The correspondence with the goals for OSI routeing stated in ISO/IEC TR 9575 are noted.

- **Network Layer Protocol Compatibility:** It is compatible with ISO 8473 and ISO 9542. (See 7.5 of ISO/IEC TR 9575),
- **Simple End systems:** It requires no changes to End systems, nor any functions beyond those supplied by ISO 8473 and ISO 9542. (See 7.2.1 of ISO/IEC TR 9575),
- **Multiple Organisations:** It allows for multiple routeing and administrative domains through the provision of static routeing information at domain boundaries. (See 7.3 of ISO/IEC TR 9575),

- **Deliverability:** It accepts and delivers NPDUs addressed to reachable destinations and rejects NPDUs addressed to destinations known to be unreachable,

- **Adaptability:** It adapts to topological changes within the routeing domain, but not to traffic changes, except potentially as indicated by local queue lengths. It splits traffic load on multiple equivalent paths. (See 7.7 of ISO/IEC TR 9575),

- **Promptness:** The period of adaptation to topological changes in the domain is a reasonable function of the domain diameter (that is, the maximum logical distance between End Systems within the domain) and Data link speeds. (See 7.4 of ISO/IEC TR 9575),

- **Efficiency:** It is both processing and memory efficient. It does not create excessive routeing traffic overhead. (See 7.4 of ISO/IEC TR 9575),

- **Robustness:** It recovers from transient errors such as lost or temporarily incorrect routeing PDUs. It tolerates imprecise parameter settings. (See 7.7 of ISO/IEC TR 9575),

- **Stability:** It stabilises in finite time to "good routes", provided no continuous topological changes or continuous data base corruptions occur,

- **System Management control:** System Management can control many routeing functions via parameter changes, and inspect parameters, counters, and routes. It will not, however, depend on system management action for correct behaviour,

- **Simplicity:** It is sufficiently simple to permit performance tuning and failure isolation,

- **Maintainability:** It provides mechanisms to detect, isolate, and repair most common errors that may affect the routeing computation and data bases. (See 7.8 of ISO/IEC TR 9575),

- **Heterogeneity:** It operates over a mixture of network and system types, communication technologies, and topologies. It is capable of running over a wide variety of subnetworks, including, but not limited to: ISO 8802 LANs, ISO/IEC 8208 and X.25 subnetworks, PSTN networks, and the OSI Data Link Service. (See 7.1 of ISO/IEC TR 9575),

- **Extensibility:** It accommodates increased routeing functions, leaving earlier functions as a subset,

- **Evolution:** It allows orderly transition from algorithm to algorithm without shutting down an entire domain,

- **Deadlock Prevention:** The congestion control component prevents buffer deadlock,

- **Very Large Domains:** With hierarchical routeing, and a very large address space, domains of essentially unlimited size can be supported. (See 7.2 of ISO/IEC TR 9575),

- **Area Partition Repair:** It permits the utilisation of level 2 paths to repair areas which become partitioned due to failing level 1 links or ISs. (See 7.7 of ISO/IEC TR 9575),

- **Determinism:** Routes are a function only of the physical topology, and not of history. In other words, the same topology will always converge to the same set of routes,

- **Protection from Mis-delivery:** The probability of mis-delivering a NPDU, i.e. delivering it to a Transport entity in the wrong End System, is extremely low,

- **Availability:** For domain topologies with cut set greater than one, no single point of failure will partition the domain. (See 7.7 of ISO/IEC TR 9575),

- **Service Classes:** The service classes of *transit delay, expense*[1], and *residual error probability* of ISO 8473 are supported through the optional inclusion of multiple routeing metrics,

- **Authentication:** The protocol is capable of carrying information to be used for the authentication of Intermediate systems in order to increase the security and robustness of a routeing domain. The specific mechanism supported in this International Standard however, only supports a weak form of authentication using passwords, and thus is useful only for protection against accidental misconfiguration errors and does not protect against any serious security threat. In the future, the algorithms may be enhanced to provide stronger forms of authentication than can be provided with passwords without needing to change the PDU encoding or the protocol exchange machinery.

6.6.2 Non-goals

The following are not within the design scope of the intra-domain IS–IS routeing protocol described in this International Standard:

- **Traffic adaptation:** It does not automatically modify routes based on global traffic load,

- **Source-destination routeing:** It does not determine routes by source as well as destination,

- **Guaranteed delivery:** It does not guarantee delivery of all offered NPDUs,

- **Level 2 Subdomain partition repair:** It will not utilise Level 1 paths to repair a level 2 subdomain partition. For full logical connectivity to be available, a connected level 2 subdomain is required,

- **Equal treatment for all ES implementations:** The End system poll function defined in 8.4.5 presumes that End systems have implemented the Suggested ES Configuration Timer option of ISO 9542. An End system which does not implement this option may experience a temporary loss of connectivity following certain types of topology changes on its local subnetwork.

1) "Expense" is referred to as "cost" in ISO 8473. The latter term is not used here because of possible confusion with the more general usage of the term to indicate path cost according to any routeing metric.

6.7 Environmental requirements

For correct operation of the protocol, certain guarantees are required from the local environment and the Data Link Layer.

6.7.1 The required local environment guarantees are:

a) Resource allocation such that the certain minimum resource guarantees can be met, including

 1) memory (for code, data, and buffers)

 2) processing;

 See 12.2.4 for specific performance levels required for conformance

b) A quota of buffers sufficient to perform routeing functions;

c) Access to a timer or notification of specific timer expiration; and

d) A very low probability of corrupting data.

6.7.2 The required subnetwork guarantees for point-to-point links are:

a) Provision that both source and destination systems complete start-up before PDU exchange can occur;

b) Detection of remote start-up;

c) Provision that no old PDUs be received after start-up is complete;

d) Provision that no PDUs transmitted after a particular startup is complete are delivered out of sequence;

e) Provision that failure to deliver a specific subnetwork SDU will result in the timely disconnexion of the subnetwork connection in both directions and that this failure will be reported to both systems; and

f) Reporting of other subnetwork failures and degraded subnetwork conditions.

g) The following events are "very low probability", which means that performance will be impacted unless they are extremely rare, on the order of less that one event per four years

 1) Delivery of NPDUs with undetected data corruption.

6.7.3 The required subnetwork guarantees for broadcast links are:

a) Multicast capability, i.e., the ability to address a subset of all connected systems with a single PDU;

b) The following events are "low probability", which means that they occur sufficiently rarely so as not to impact performance, on the order of once per thousand PDUs

 1) Routeing PDU non-sequentiality,

 2) Routeing PDU loss due to detected corruption; and

 3) Receiver overrun;

c) The following events are "very low probability", which means performance will be impacted unless they are extremely rare, on the order of less than one event per four years

 1) Delivery of NPDUs with undetected data corruption; and

 2) Non-transitive connectivity, i.e. where system *A* can receive transmissions from systems *B* and *C*, but system *B* cannot receive transmissions from system *C*.

6.7.4 The following services are assumed to be not available from broadcast links:

a) Reporting of failures and degraded subnetwork conditions that result in NPDU loss, for instance receiver failure. The routeing functions are designed to account for these failures.

6.8 Functional organisation of subnetwork independent components

The Subnetwork Independent Functions are broken down into more specific functional components. These are described briefly in this sub-clause and in detail in clause 7. This International Standard uses a functional decomposition adapted from the model of routeing presented in subclause 5.1 of ISO/IEC TR 9575. The decomposition is not identical to that in ISO/IEC TR 9575, since that model is more general and not specifically oriented toward a detailed description of intra-domain routeing functions such as supplied by this protocol.

The functional decomposition is shown below in figure 3.

The routeing processes are:

- *Decision Process*
- *Update Process*

 NOTE 4 This comprises both the *Information Collection* and *Information Distribution* components identified in ISO/IEC TR 9575.

- *Forwarding Process*
- *Receive Process*

6.8.1 Decision process

This process calculates routes to each destination in the domain. It is executed separately for level 1 and level 2 routeing, and separately within each level for each of the routeing metrics supported by the Intermediate system. It uses the *Link State Database*, which consists of information from the latest Link State PDUs from every other Intermediate system in the area, to compute shortest paths from this IS to all other systems in the area – ⑨ in figure 3. The Link State Data Base is maintained by the Update Process.

Execution of the Decision Process results in the determination of [circuit, neighbour] pairs (known as *adjacencies*), which are stored in the appropriate Forwarding Information base – ⑩ – and used by the Forwarding process as paths along which to forward NPDUs.

Several of the parameters in the routeing data base that the Decision Process uses are determined by the implementation. These include:

- maximum number of Intermediate and End systems within the IS's area;

- maximum number of Intermediate and End system neighbours of the IS, etc.,

so that databases can be sized appropriately. Also parameters such as

- routeing metrics for each circuit; and

- timers;

can be adjusted for enhanced performance. The complete list of System Management set-able parameters is contained in clause 11.

6.8.2 Update process

This process constructs, receives and propagates Link State PDUs. Each Link State PDU contains information about the identity and routeing metric values of the adjacencies of the IS that originated the Link State PDU.

The Update Process receives Link State and Sequence Numbers PDUs from the Receive Process — ④ in figure 3. It places new routeing information in the routeing information base — ⑥ and propagates routeing information to other Intermediate systems – ⑦ and ⑧ .

General characteristics of the Update Process are:

- Link State PDUs are generated as a result of topological changes, and also periodically. They may also be generated indirectly as a result of System Management actions (such as changing one of the routeing metrics for a circuit).

- Level 1 Link State PDUs are propagated to all Intermediate systems within an area, but are not propagated out of an area.

- Level 2 Link State PDUs are propagated to all Level 2 Intermediate systems in the domain.

- Link State PDUs are not propagated outside of a domain.

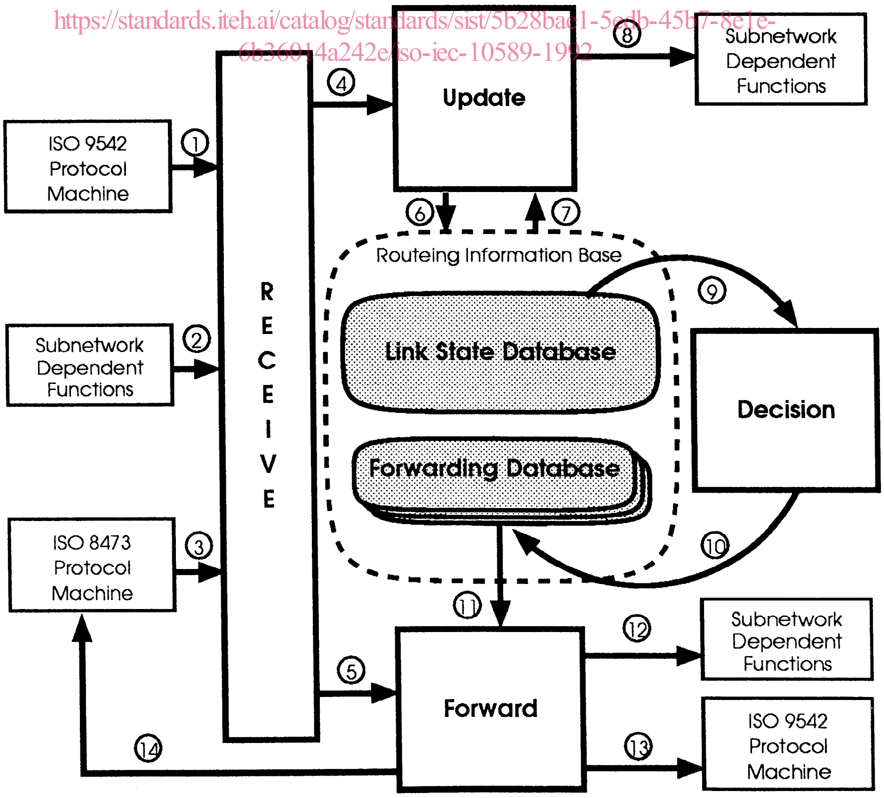

Figure 3 - Decomposition of Subnetwork Independent Functions

- The update process, through a set of System Management parameters, enforces an upper bound on the amount of routeing traffic overhead it generates.

6.8.3 Forwarding process

This process supplies and manages the buffers necessary to support NPDU relaying to all destinations.

It receives, via the Receive Process, ISO 8473 PDUs to be forwarded – (5) in figure 3.

It performs a lookup in the appropriate[1)] Forwarding Database – (11) – to determine the possible output adjacencies to use for forwarding to a given destination, chooses one adjacency – (12) —, generates error indications to ISO 8473 – (14) , and signals ISO 9542 to issue Redirect PDUs – (13).

6.8.4 Receive process

The Receive Process obtains its inputs from the following sources

- received PDUs with the NLPID of Intra-Domain routeing – (2) in figure 3,
- routeing information derived by the ES–IS protocol from the receipt of ISO 9542 PDUs – (1) ; and
- ISO 8473 data PDUs handed to the routeing function by the ISO 8473 protocol machine – (3) .

It then performs the appropriate actions, which may involve passing the PDU to some other function (e.g. to the Forwarding Process for forwarding – (5)).

7 Subnetwork independent functions

This clause describes the algorithms and associated databases used by the routeing functions. The managed objects and attributes defined for System Management purposes are described in clause 11.

The following processes and data bases are used internally by the subnetwork independent functions. Following each process or data base title, in parentheses, is the type of systems which must keep the database. The system types are “L2” (level 2 Intermediate system), and “L1” (level 1 Intermediate system). Note that a level 2 Intermediate system is also a level 1 Intermediate system in its home area, so it must keep level 1 databases as well as level 2 databases.

Processes:

- Decision Process (L2, L1)
- Update Process (L2, L1)
- Forwarding Process (L2, L1)
- Receive Process (L2, L1)

Databases:

- Level 1 Link State data base (L2, L1)
- Level 2 Link State data base (L2)
- Adjacency Database (L2, L1)
- Circuit Database (L2, L1)
- Level 1 Shortest Paths Database (L2, L1)
- Level 2 Shortest Paths Database (L2)
- Level 1 Forwarding Databases — one per routeing metric (L2, L1)
- Level 2 Forwarding Database — one per routeing metric (L2)

7.1 Addresses

The NSAP addresses and NETs of systems are variable length quantities that conform to the requirements of ISO 8348/Add.2. The corresponding NPAI contained in ISO 8473 PDUs and in this protocol's PDUs shall use the preferred binary encoding. Any of the AFIs and their corresponding DSP syntax may be used with this protocol.

7.1.1 Address structure for intradomain IS–IS routeing

In order to understand the requirements set under the present clause 7.1, it is necessary to view the encoded NSAPs or NETs as structured according to figure 4, where three fields are distinguished:

a) Area Address
b) ID
c) SEL

1) The appropriate forwarding database is selected by choosing a routeing metric based on fields in the QoS Maintenance option field of ISO 8473.

7.1.2 NPAI — area address field

An area address is a variable length quantity consisting of the entire high-order part of the NPAI, excluding the ID and SEL fields.

7.1.3 NPAI of systems within a routeing domain

The structure of the ID and SEL fields of the NPAI are interpreted in the following way by the protocol defined in this International Standard:

ID System identifier — a variable length field from 1 to 8 octets (inclusive). Each routeing domain employing this protocol shall select a single size for the ID field and all Intermediate systems in the routeing domain shall use this length for the system IDs of all systems in the routeing domain.

The set of ID lengths supported by an implementation is an implementation choice, provided that at least one value in the permitted range can be accepted. The routeing domain administrator must ensure that all ISs included in a routeing domain are able to use the ID length chosen for that domain.

SEL NSAP Selector — a 1-octet field which acts as a selector for the entity which is to receive the PDU (this may be a Transport entity or the Intermediate system Network entity itself). It is the least significant (last) octet of the NPAI.

7.1.4 Administration and deployment of systems in a routeing domain

It is the responsibility of the routeing domain administrative authority to enforce the requirements stated below in this clause. These requirements place specific constraints on the NSAP addresses and NETs of systems deployed in a routeing domain, when these systems operate the protocol defined in this International Standard. The protocol defined in this International Standard assumes that these requirements are met, but has no means to verify compliance with them.

NOTE 5 To correctly interpret the requirements given below, it is necessary to refer both to the structure of the NPAI presented in 7.1.1, and to the concept of manual area addresses defined in 7.1.5.

For correct operation of the routeing protocol defined in this International Standard, the following requirements must be met in a routeing domain:

a) For all systems in the routeing domain:

   1) By definition, all systems in a routeing domain that have a given value of area address belong to the same area.

      NOTE 6 A consequence of this requirement is that a reachable address prefix may not match any area address of an area in the routeing domain. However, an IS is not required to perform any dynamic check to detect if this property is violated due to system management misconfiguration.

   2) Each system in an area must have an unambiguous ID; that is, no two systems (IS or ES) in an area may use the same ID value.

   3) All systems belonging to a given routeing domain must have NETs or NSAP addresses whose ID fields are of equal length.

b) Additional requirements for Intermediate system addresses:

   1) Each Level 2 Intermediate system within a routeing domain must have an unambiguous value for its ID field; that is, no two level 2 ISs in a routeing domain can have the same value in their ID fields.

c) Additional requirements for area administration:

   1) No two End systems in an area may have addresses that match in all but the SEL field.

d) Requirements placed on End systems to be neighbours of a level 1 IS:

   An End system may be a neighbour of a level 1 IS if and only if:

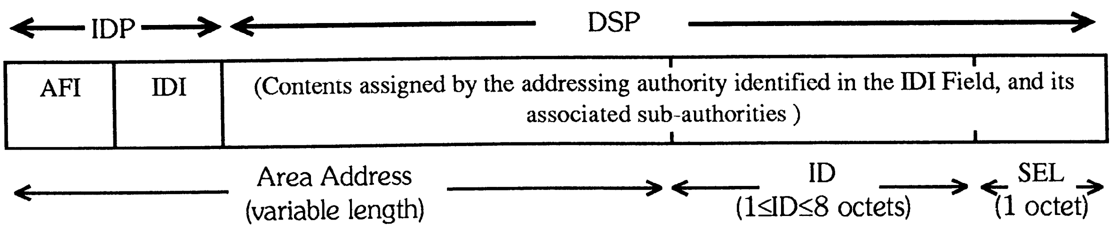

NOTE - The diagram shows how this International Standard will parse the NPAI (that is, an address or NET encoded by the preferred binary encoding of ISO 8348/Add.2). This protocol uses the fields "Area Address", "ID", and "SEL".

To emphasize that address information is constructed according to ISO 8348/Add.2, the quantities "IDP", "DSP", "AFI", and "IDI" are also shown; however, they are not explicitly used by this protocol.

Figure 4 - Address Structure for Intra-domain IS-IS Routeing