
**Information technology —
Telecommunications and information
exchange between systems — Intermediate
System to Intermediate System
intra-domain routing 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)*



PDF disclaimer

This PDF file may contain embedded typefaces. In accordance with Adobe's licensing policy, this file may be printed or viewed but shall not be edited unless the typefaces which are embedded are licensed to and installed on the computer performing the editing. In downloading this file, parties accept therein the responsibility of not infringing Adobe's licensing policy. The ISO Central Secretariat accepts no liability in this area.

Adobe is a trademark of Adobe Systems Incorporated.

Details of the software products used to create this PDF file can be found in the General Info relative to the file; the PDF-creation parameters were optimized for printing. Every care has been taken to ensure that the file is suitable for use by ISO member bodies. In the unlikely event that a problem relating to it is found, please inform the Central Secretariat at the address given below.

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:2002

© ISO/IEC 2002

All rights reserved. Unless otherwise specified, no part of this publication may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying and microfilm, without permission in writing from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
Case postale 56 • CH-1211 Geneva 20
Tel. + 41 22 749 01 11
Fax + 41 22 749 09 47
E-mail copyright@iso.org
Web www.iso.org

Published in Switzerland

Contents

1	<i>Scope</i>	1
2	<i>Normative references</i>	1
3	<i>Definitions</i>	3
3.1	Reference model definitions	3
3.2	Network layer architecture definitions	3
3.3	Network layer addressing definitions	3
3.4	Local area network definitions	3
3.5	Routeing framework definitions	3
3.6	Additional definitions	3
4	<i>Symbols and abbreviations</i>	5
4.1	Data units	5
4.2	Protocol data units	5
4.3	Addresses	5
4.4	Miscellaneous	5
5	<i>Typographical conventions</i>	6
6	<i>Overview of the protocol</i>	6
6.1	System types	6
6.2	Subnetwork types	7
6.3	Topologies	7
6.4	Addresses	8
6.5	Functional organisation	8
6.6	Design goals and non-goals	9
6.7	Environmental requirements	11
6.8	Functional organisation of subnetwork independent components	12
7	<i>Subnetwork independent functions</i>	14
7.1	Addresses	15
7.2	Decision process	18
7.3	Update process	26
7.4	Forwarding process	45
7.5	Routeing constants and parameters	48
8	<i>Subnetwork dependent functions</i>	49
8.1	<i>Multi-destination circuits on ISs at a domain boundary</i>	49
8.2	Point-to-point subnetworks	50
8.3	ISO 8208 subnetworks	54
8.4	Broadcast subnetworks	59

9	<i>Structure and encoding of PDUs</i>	65
9.1	General encoding rules	65
9.2	Encoding of network layer addresses	65
9.3	Encoding of SNPA addresses	65
9.4	PDU types	66
9.5	Level 1 LAN IS to IS hello PDU	66
9.6	Level 2 LAN IS to IS hello PDU	69
9.7	Point-to-point IS to IS hello PDU	72
9.8	Level 1 link state PDU	75
9.9	Level 2 link state PDU	79
9.10	Level 1 complete sequence numbers PDU	84
9.11	Level 2 complete sequence numbers PDU	86
9.12	Level 1 partial sequence numbers PDU	88
9.13	Level 2 partial sequence numbers PDU	90
10	<i>System environment</i>	91
10.1	Generating jitter on timers	91
10.2	Resolution of timers	92
10.3	Requirements on the operation of ISO 9542	93
10.4	Requirements on the operation of ISO 8473	93
11	<i>System management</i>	93
11.1	General	93
--11.2	<i>GDMO definition</i>	94
--11.2.1	Common GDMO definitions	94
--11.3	ASN1 modules	127
12	<i>Conformance</i>	129
12.1	Conformance for protocol implementation	129
12.1.2	Dynamic conformance	131
12.2	Conformance for management information implementation	133
Annex A		135
A.1	Introduction	135
A.2	Abbreviations and special symbols	135
A.3	Instructions for completing the pics pro formas	135
A.4	Identification	138
A.5	Protocol summary: ISO 10589 general	139
A.6	Protocol summary: ISO 10589 level 1 specific functions	143
A.7	Protocol summary: ISO 10589 level 2 specific functions	144

Annex B	146
B.1 Addressing and routeing	146
B.2 Use of the area address field in intra-domain routeing	148
Annex C	150
C.1 Routeing databases	150
C.3 Forwarding process	156
Annex D	158
D.1 Congestion control	158
D.2 Congestion avoidance	159
Annex E	160
--E.1 Generic managed object class definitions	160
--E.2 ASN.1 definitions	167
Annex F	169
Annex G	170
G.1 Introduction	170
G.2 Identification of the implementation	171
G.3 Identification of the International Standard in which the management information is defined	171
Annex H	175
H.1 Introduction	175
H.2 Instructions for completing the MICS proforma to produce a MICS	175
H.3 Symbols, abbreviates and terms	175
H.4 Statement of Conformance to the management information	175
Annex I	182
I.1 Introduction	182
I.2 Adjacency managed object	182
I.3 Virtual adjacency managed object	185
I.4 Destination system managed object	187
I.5 Destination area managed object	188
I.6 reachableAddress ["ISO/IEC 10589"]	190
Annex J	195
J.1 Introduction	195
J.2 Instructions for completing the MRCS proforma for name binding to produce a MRCS	195
J.3 Statement of conformance to the name binding	195
Index	197

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.

International Standards are drafted in accordance with the rules given in the ISO/IEC Directives, Part 3.

The main task of the joint technical committee is to prepare International Standards. 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.

Attention is drawn to the possibility that some of the elements of this International Standard may be the subject of patent rights. ISO and IEC shall not be held responsible for identifying any or all such patent rights.

ISO/IEC 10589 was prepared by Joint Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 6, *Telecommunications and information exchange between systems*.

This second edition cancels and replaces the first edition (ISO/IEC 10589:1992), which has been technically revised. It incorporates Cor.1:1993, Cor.2:1996, Cor.3:1996, Amd.1:1996 and Amd.2:1999.

Annexes A, E, G, H, I and J form a normative part of this International Standard. Annexes B, C, D and F are for information only.

Annexes G, H, I and J provide ICS proformas associated with intra-domain routing protocol management information.

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 routing domain to exchange configuration and routing information to facilitate the operation of the routing 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 routing protocol is intended to support large routing 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 routing domains, provision is made for intra-domain routing to be organised hierarchically. A large domain may be administratively divided into *areas*. Each system resides in exactly one area. Routing within an area is referred to as *Level 1 routing*. Routing between areas is referred to as *Level 2 routing*. Level 2 Intermediate Systems keep track of the paths to destination areas. Level 1 Intermediate Systems keep track of the routing 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 routing to the destination area, where it again travels via level 1 routing to the destination End System.

Information technology – Telecommunications and information exchange between systems – Intermediate System to Intermediate System intra-domain routing 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 routing information for the purpose of routing within a single routing 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 routing information between network entities residing in Intermediate Systems within a single routing domain;
- b) the encoding of the protocol data units used for the transmission of the configuration and routing 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

- a) the interactions between Intermediate system Network entities through the exchange of protocol data units;
- b) the interactions between a Network entity and an underlying service provider through the exchange of subnetwork service primitives; and
- c) the constraints on route determination which must be observed by each Intermediate system when each has a routing information base which is consistent with the others.

2 Normative references

The following normative documents contain provisions which, through reference in this text, constitute provisions of this International Standard. For dated references, subsequent amendments to, or revisions of, any of these publications do not apply. However, parties to agreements based on this International Standard are encouraged to investigate the possibility of applying the most recent editions of the normative documents indicated below. For undated references, the latest edition of the normative document referred to applies. Members of ISO and IEC maintain registers of currently valid International Standards.

ISO/IEC 7498-1:1994, *Information technology – Open Systems Interconnection – Basic Reference Model: The Basic Model*

ISO/IEC 7498-3:1997, *Information technology – Open Systems Interconnection – Basic Reference Model: Naming and addressing*

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

ISO/IEC 8208:2000, *Information technology – Data communications – X.25 Packet Layer Protocol for Data Terminal Equipment*

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

ISO/IEC 8348:1996, *Information technology – Open Systems Interconnection – Network Service Definition*

ISO/IEC 8473-1:1998, *Information technology – Protocol for providing the connectionless-mode network service: Protocol specification*

ISO/IEC 8473-4:1995, *Information technology – Protocol for providing the connectionless-mode network service: Provision of the underlying service by a subnetwork that provides the OSI data link service*

ISO 8648:1988, *Information processing systems – Open Systems Interconnection – Internal organization of the Network Layer*

ISO/IEC TR 8802-1:1997, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 1: Overview of Local Area Network Standards*

ISO/IEC 8802-2:1998, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 2: Logical link control*

ISO/IEC 8802-3:2000, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*

ISO/IEC 8802-5:1998, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 5: Token ring access method and physical layer specifications*

ISO/IEC 8802-6:1994, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 6: Distributed Queue Dual Bus (DQDB) access method and physical layer specifications*

ISO/IEC 9314 (all parts), *Information technology – Fibre Distributed Data Interface (FDDI)*

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

ISO/IEC TR 9575:1995, *Information technology – Telecommunications and information exchange between systems – OSI Routing Framework*

ISO/IEC TR 9577:1999, *Information technology – Protocol identification in the network layer*

ISO/IEC 15802-1:1995, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Common specifications – Part 1: Medium Access Control (MAC) service definition*

ISO/IEC 10165-1:1993, *Information technology – Open Systems Interconnection – Management Information Services – Structure of management information: Management Information Model*

ISO/IEC 10165-4:1992, *Information technology – Open Systems Interconnection – Structure of management information – Part 4: Guidelines for the definition of managed objects*

ISO/IEC 10733:1998, *Information technology – Elements of management information related to the OSI Network Layer*

ISO/IEC 8824-1:1998, *Information technology – Abstract Syntax Notation One (ASN.1): Specification of basic notation*

ISO/IEC 8825-1:1998, *Information technology – ASN.1 encoding rules: Specification of Basic Encoding Rules (BER), Canonical Encoding Rules (CER) and Distinguished Encoding Rules (DER)*

ISO/IEC 9646-7:1995, *Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 7: Implementation Conformance Statements*

ISO/IEC 10165-6:1997, *Information technology – Open Systems Interconnection – Structure of management information: Requirements and guidelines for implementation conformance statement proformas associated with OSI management*

NOTE 1 – ISO/IEC 9646-1:1994 and ISO/IEC 9646-2:1994 supersede ISO/IEC 9646-1:1991 and ISO/IEC 9646-2:1991 respectively. However, when this International Standard was under development, the previous editions were valid and this International Standard is therefore based on these editions, which are listed below.

ISO/IEC 9646-1:1991, *Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 1: General concepts*

ISO/IEC 9646-2:1991, *Information technology – Open Systems Interconnection – Conformance testing methodology and framework – Part 2: Abstract test suite specification*

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:

- 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 routing subdomain which maintains detailed routing information about its own internal composition, and also maintains routing information which allows it to reach other routing 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 routing 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 routing domain.

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

3.6.4 circuit: A subset of the local routing 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 routing subdomain: a set of Intermediate systems and End systems located within the same Routing domain.

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




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

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 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

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
MCS	Management conformance summary
MICS	Management information conformance statement
MOCS	Managed object conformance statement
MRCS	Managed relationship conformance statement
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 Point
SNDP	Subnetwork Dependent Convergence Protocol
SNICP	Subnetwork Independent Convergence Protocol
SRM	Send Routing Message
SSN	Send Sequence Numbers
SVC	Switched Virtual Circuit




Figure 2 Topologies and Systems supported by Intradomain Routing

5 Typographical conventions

This International Standard makes use of the following typographical conventions:

- important terms and concepts appear in *italic* type when introduced for the first time;
- 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;
- protocol field names appear in **sansSerif** type with each word capitalised; and
- 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 destination 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 routing domain.

NOTE 2 Operation of a level 2 IS may be restricted on a per circuit basis by setting `manualL2OnlyMode` to "True". This indicates that the circuit is to be used only for Level 2 traffic. If all circuits in a Level 2 IS have `manualL2OnlyMode` set to "True" then the IS does not operate as a Level 1 IS in the area.

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 sub-networks:

- 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 routing PDUs are exchanged between ISs on a DA circuit.

NOTE 3 For the operation of the protocol, a permanent point-to-point link and a static point-to-point DED are equivalent (see 7.3.7, 7.3.9).

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 *Routing Domains*. This has certain advantages, as described in ISO/IEC TR 9575. Furthermore, it is desirable for an intra-domain routing protocol to aid in the operation of an inter-domain routing protocol, where such a protocol exists for interconnecting multiple routing domains.

In order to facilitate the construction of such multi-domain topologies, provision is made for the entering of inter-domain routing 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 routing protocol at the ISs which have links crossing routing 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 routing 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 routing) by that IS and delivered out of the domain. (It is assumed that the routing functions of the other domain will then be able to deliver the NPDu to its destination.)

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

NOTE 4 This model also permits the construction of routing domains whose scope is not limited by the hierarchical nature of network layer address assignment. For example, it is possible to construct a routing domain, or even a single area, whose area addresses are taken from multiple addressing authorities.

6.4 Addresses

Within a routing 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 routing 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. However, the routing 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, but these End systems may require additional configuration information to be entered into the Intermediate systems and they may obtain inferior routing performance.

NOTE 5 The procedures whereby the routing 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 routing 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 6 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:

- **Routing.** The routing 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 routing domain is used to ascertain the existence of a path, and route the NPDU to its destination. The routing component at an Intermediate system has the following specific functions:

- It extracts and interprets the routing 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, 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 routing 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 routing and administrative domains through the provision of static routing 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 routing 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 routing traffic overhead. (See 7.4 of ISO/IEC TR 9575),
- **Robustness:** It recovers from transient errors such as lost or temporarily incorrect routing 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 routing 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 routing 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 routing 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 routing, 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 Misdelivery:** The probability of misdelivering 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 routing 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 routing 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.

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 routing metric.

6.6.2 Non-goals

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

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

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 routing 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 disconnection 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 sub-network 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 than 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) Routing PDU non-sequentiality,

- 2) Routing 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:

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

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 routing 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 routing functions such as supplied by this protocol.

The functional decomposition is shown below in figure 3.

The routing processes are:

- *Decision Process*
- *Update Process*

NOTE 7 This comprises both the *Information Collection* and *Information Distribution* components 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 routing, and separately within each level for each of the routing 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 routing 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

- routing metrics for each circuit; and
- timers; can be adjusted for enhanced performance. The complete list of System Management settable 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 routing 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 routing information in the routing information base — 6 and propagates routing 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 routing 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.
- The update process, through a set of System Management parameters, enforces an upper bound on the amount of routing 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 appropriate¹⁾ 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 routing — 2 in figure 3,
- routing information derived by the ES-IS protocol from the receipt of ISO 9542 PDUs — 1 —; and
- ISO 8473 data PDUs handed to the routing 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).

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




Figure 3 - Decomposition of Subnetwork Independent Functions

7 Subnetwork independent functions

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

The following processes and databases 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 routing metric (L2, L1)
- Level 2 Forwarding Database — one per routing metric (L2)

7.1 Addresses

The NSAP addresses and NETs of systems are variable length quantities that conform to the requirements of ISO 8348. 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 routing

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

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 routing domain

This clause first defines how the NPAI corresponding to NSAP addresses and Network Entity Titles of systems deployed in a routing domain is constructed and second how the NPAI is structured for use by the protocol.

7.1.3.1 Construction of NPAI from network addresses

NPAI is derived from NSAP addresses and NETs according to ISO/IEC 8348. The NETs and NSAP addresses are obtained from the appropriate addressing authorities.

For these addresses to be correctly interpreted by the protocol in this International Standard it is mandatory that the routing domain authority ensure that whenever the Network address includes a DSP whose syntax is decimal digits, then:

- if the IDI is in the ISO DCC or E.164 format, the decimal syntax DSP must be an odd number of digits;
- for the other IDI formats the decimal syntax DSP must be an even number of decimal digits.

7.1.3.2 Structure of the NPAI

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 routing domain employing this protocol shall select a single size for the ID field and all Intermediate systems in the routing domain shall use this length for the system IDs of all systems in the routing 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 routing domain administrator must ensure that all ISs included in a routing 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.

NOTE 8 The SEL field is always the last octet of the NPAI, since the rules enforced in 7.1.3.1 guarantee that there will be no pad at the end of the NPAI.

7.1.4 Administration and deployment of systems in a routing domain

It is the responsibility of the routing domain administrative authority to enforce the requirements stated below in this clause and the rules stated above in 7.1.3.1. These requirements place specific constraints on the NSAP addresses and NETs of systems deployed in a routing 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 9 To interpret correctly the requirements given below, it is necessary to refer to the structure of the NPAI presented in 7.1.1 and 7.1.3.2 and to the concept of manual area addresses defined in 7.1.5.

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

a) For all systems in the routing domain:

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

NOTE 10 A consequence of this requirement is that a reachable address prefix may not match any area address of an area in the routing 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 routing 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 routing domain must have an unambiguous value for its ID field; that is, no two level 2 ISs in a routing 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:


- 1) its area address matches the level 1 IS's area address as contained in its own NET, which is always an entry in the adjacent IS's `manualAreaAddresses` parameter, or
- 2) its area address matches one of the other entries in the adjacent IS's `manualAreaAddresses` parameter.

NOTE 11 The definitions of several entries in the adjacent IS's `manualAreaAddresses` parameter allow for systems with differing area addresses to be merged under one level 1 IS for the purpose of fabricating a single area. This concept is described in detail under 7.1.5 below.

e) Additional requirements for Interdomain routing:

When Interdomain routing between two routing domains is assisted through the use of reachable address prefixes, then

- 1) no address of any area in one of the routing domains should match the address of an area in the other routing domain.



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). 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 Routing

7.1.5 Manual area addresses

Within a routing domain, it is often convenient to associate more than one area address with an area. There are a number of reasons why assigning more than one area address may be useful, including the following.

- a) There may be more than one addressing authority involved in the assignment of addresses in the routing domain, yet it is not efficient to require a separate area for each addressing domain.
- b) At times it may be necessary to reconfigure a routing domain by dividing an area into two or more areas, or combining a number of areas into a single area. These reconfigurations could not be done during normal routing domain operation if only a single area address per area were permitted.

Therefore, this International Standard permits an area to have a number of synonymous area addresses.

The use of several synonymous area addresses by an IS is accommodated through the use of the management parameter **manualAreaAddresses**. This parameter is set locally for each IS by system management; it contains a list of all synonymous area addresses associated with the IS. All of the IS's **manualAreaAddresses**, when combined with the IS's **systemID**, are valid network entity titles for the IS.

Each level 1 IS distributes its **manualAreaAddresses** in its Level 1 LSP's **Area Addresses** field, thus allowing level 2 ISs to create a composite list of all area addresses in use within a given area. Level 2 ISs in turn advertise the composite list throughout the level 2 subdomain by including it in their Level 2 LSP's **Area Addresses** field, thus distributing information on all the area addresses associated with the entire routing domain. The procedures for establishing an adjacency between two level 1 ISs require that there be at least one area address in common between their two **manualAreaAddresses** lists, and the procedures for establishing an adjacency between a level 1 IS and an End system require that the End system's area address match an entry in the IS's **manualAreaAddresses** list. Therefore, it is the responsibility of System Management to ensure that each area address associated with an IS is included. In particular, system management must ensure that the area addresses of all ESs and Level 1 ISs adjacent to a given level 1 IS are included in that IS's **manualAreaAddresses** list.

The union of all area addresses of the ISs in an area may exceed the configured capacity of one or more ISs in the area. In order that all ISs agree on the area addresses of an area, each IS has a parameter **maximumAreaAddresses**, established by System Management. The value of **maximumAreaAddresses** is communicated among the ISs in the protocol's PDUs and is checked to ensure that all ISs in an area have the same value of this parameter. All ISs shall support a value of at least 3 for **maximumAreaAddresses**, although the value may be set lower by System Management if desired. Failure to set **maximumAreaAddresses** consistently among the ISs in an area may cause adjacencies to fail to initialise and/or an area to become partitioned.

If the area address field for the destination address of an ISO 8473 PDU — or for the next entry in its source routing field, when present — is not listed in the parameter **areaAddresses** of a level 1 IS receiving the PDU, then the destination system does not reside in the IS's area. Such PDUs will be routed by level 2 routing.

7.1.6 Encoding of addressing information

This International Standard makes use of four types of address information: NETs, NSAP addresses, area addresses, and address prefixes. The encoding rules for each of them are given below.

- a) NETs shall be encoded according to the preferred binary encoding specified in ISO 8348.
- b) NSAP addresses shall be encoded according to the preferred binary encoding specified in ISO 8348.
- c) The encoded form of an area address shall be obtained by dropping the last $IDLength + 1$ octets of the preferred binary encoding of the corresponding NSAP, where $IDLength$ is equal to the length of the ID field used by the routing domain.
- d) The encoded form of an address prefix shall be obtained by encoding the prefix (expressed in its abstract syntax), according to the preferred binary encoding, unless the end of the prefix falls within the IDP. In this case, each decimal digit in the prefix shall be encoded as the corresponding semi-octet in the range 0000-1001 and no padding characters shall be inserted.

7.1.7 Matching an NSAP address with an area address or an address prefix

A destination NSAP address can be matched against either an area address or an address prefix. An area address or an address prefix which extends into the DSP, shall be compared directly against the encoded NSAP address, including any padding

characters that may be present; an address prefix which does not extend into the DSP shall be compared against NSAP', which is obtained from the encoded NSAP address by removing all padding characters that were inserted by the binary encoding.

The existence of a match shall be determined as follows:

- a) If the encoded NSAP (or NSAP') contains fewer semi-octets than the encoded area address (or address prefix), then there is no match.
- b) If the NSAP (or NSAP') contains at least as many octets as the area address (or address prefix), and all octets of the encoded area address (or address prefix) are identical to the corresponding leading octets of the encoded NSAP address (or NSAP'), there is a match. Otherwise, there is no match.

NOTE 12 Any implementation of a matching process that satisfies the requirements listed above may be used. The key point is that matching process must be aware of whether or not the encoded area address or address prefix extends into the DSP, and must then either include or exclude padding characters from the encoded NSAP, as defined above.

7.1.8 Comparison of addresses

Unless otherwise stated, comparison of addresses shall be performed on the addresses encoded as sequences of octets in the form specified in 9.2 or 9.3 as appropriate.

If the encoded addresses are of different lengths and the shorter encoded address is a prefix of the longer, the shorter address is considered to be less than the longer.

Otherwise, the result of the comparison is that obtained by padding the shorter encoded address (if any) with trailing zero octets to the length of the longer, and comparing the two resulting octet sequences as unsigned integers, with the first octet of each sequence considered the most significant.

The addresses to which this procedure applies are NSAP addresses, Network Entity Titles, and SNPA addresses.

7.2 Decision process

This process uses the database of Link State information to calculate the forwarding database(s), from which the forwarding process can know the proper next hop for each NPDUs. The Level 1 Link State Database is used for calculating the Level 1 Forwarding Database(s), and the Level 2 Link State Database is used for calculating the Level 2 Forwarding Database(s).

7.2.1 Input and output

INPUT

- Link State Database – This database is a set of information from the latest Link State PDUs from all known Intermediate systems (within this area, for Level 1, or within the level 2 subdomain, for Level 2). This database is received from the Update Process.
- Notification of an Event – This is a signal from the Update Process that a change to a link has occurred somewhere in the domain.

OUTPUT

- Level 1 Forwarding Databases — one per routing metric
- (Level 2 Intermediate systems only) Level 2 Forwarding Databases — one per routing metric
- (Level 2 Intermediate systems only) The Level 1 Decision Process informs the Level 2 Update Process of the ID of the Level 2 Intermediate system within the area with lowest ID reachable with real level 1 links (as opposed to a virtual link consisting of a path through the level 2 subdomain)
- (Level 2 Intermediate systems only) If this Intermediate system is the “Partition Designated Level 2 Intermediate system” in this partition, the Level 2 Decision Process informs the Level 1 Update Process of the values of the default routing metric to and ID of the “partition designated level 2 Intermediate system” in each other partition of this area.

7.2.2 Routeing metrics

There are four routeing metrics defined, corresponding to the four possible orthogonal qualities of service defined by the QoS Maintenance field of ISO 8473. Each circuit emanating from an Intermediate system shall be assigned a value for one or more of these metrics by System management. The four metrics are as follows:

- a) *Default metric*: This is a metric understood by every Intermediate system in the domain. Each circuit shall have a positive integral value assigned for this metric. The value may be associated with any objective function of the circuit, but by convention is intended to measure the *capacity* of the circuit for handling traffic, for example, its throughput in bits-per-second. Higher values indicate a lower capacity.
- b) *Delay metric*: This metric measures the *transit delay* of the associated circuit. It is an optional metric, which if assigned to a circuit shall have a positive integral value. Higher values indicate a longer transit delay.
- c) *Expense metric*: This metric measures the *monetary cost* of utilising the associated circuit. It is an optional metric, which if assigned to a circuit shall have a positive integral value¹⁾. Higher values indicate a larger monetary expense.
- d) *Error metric*: This metric measures the *residual error probability* of the associated circuit. It is an optional metric, which if assigned to a circuit shall have a non-zero value. Higher values indicate a larger probability of undetected errors on the circuit.

NOTE 13 The decision process combines metric values by simple addition. It is important, therefore, that the values of the metrics be chosen accordingly.

Every Intermediate system shall be capable of calculating routes based on the default metric. Support of any or all of the other metrics is optional. If an Intermediate system supports the calculation of routes based on a metric, its update process may report the metric value in the LSPs for the associated circuit; otherwise, the IS shall not report the metric.

When calculating paths for one of the optional routeing metrics, the decision process only utilises LSPs with a value reported for the corresponding metric. If none of an IS's circuits has a value associated with one of the optional metrics, then IS shall not calculate routes based on that metric.

NOTE 14 A consequence of the above is that a system reachable via the default metric may not be reachable by another metric.

See 7.4.2 for a description of how the forwarding process selects one of these metrics based on the contents of the ISO 8473 QoS Maintenance option.

Each of the four metrics described above may be of two types: an *Internal metric* or an *External metric*. Internal metrics are used to describe links/routes to destinations internal to the routeing domain. External metrics are used to describe links/routes to destinations outside of the routeing domain. These two types of metrics are not directly comparable, except the internal routes are always preferred over external routes. In other words an internal route will always be selected even if an external route with lower total cost exists.

7.2.3 Broadcast subnetworks

Instead of treating a broadcast subnetwork as a fully connected topology, the broadcast subnetwork is treated as a pseudonode, with links to each attached system. Attached systems shall only report their link to the pseudonode. The designated Intermediate system, on behalf of the pseudonode, shall construct Link State PDUs reporting the links to all the systems on the broadcast subnetwork with a zero value for each supported routeing metric²⁾.

The pseudonode shall be identified by the *sourceID* of the Designated Intermediate system, followed by a non-zero *pseudonodeID* assigned by the Designated Intermediate system. The *pseudonodeID* is locally unique to the Designated Intermediate system.

Designated Intermediate systems are determined separately for level 1 and level 2. They are known as the *LAN Level 1 Designated IS* and the *LAN Level 2 Designated IS* respectively. See 8.4.5.

¹⁾ The path computation algorithm utilised in this International Standard requires that all circuits be assigned a positive value for a metric. Therefore, it is not possible to represent a "free" circuit by a zero value of the expense metric. By convention, the value 1 is used to indicate a "free" circuit.

²⁾ They are set to zero metric values since they have already been assigned metrics by the link to the pseudonode. Assigning a non-zero metric in the pseudonode LSP would have the effect of doubling the actual value.

An Intermediate system may resign as Designated Intermediate System on a broadcast circuit either because it (or its SNPA on the broadcast subnetwork) is being shut down or because some other Intermediate system of higher priority has taken over that function. When an Intermediate system resigns as Designated Intermediate System, it shall initiate a network wide purge of its pseudonode Link State PDU(s) by setting their Remaining Lifetime to zero and performing the actions described in 7.3.16.4. A LAN Level 1 Designated Intermediate System purges Level 1 Link State PDUs and a LAN Level 2 Designated Intermediate System purges Level 2 Link State PDUs. An Intermediate system which has resigned as both Level 1 and Level 2 Designated Intermediate System shall purge both sets of LSPs.

When an Intermediate system declares itself as designated Intermediate system and it is in possession of a Link State PDU of the same level issued by the previous Designated Intermediate System for that circuit (if any), it shall initiate a network wide purge of that (or those) Link State PDU(s) as above.

7.2.4 Links

Two Intermediate systems are not considered neighbours unless each reports the other as directly reachable over one of their SNPAs. On a Connection-oriented subnetwork (either point-to-point or general topology), the two Intermediate systems in question shall ascertain their neighbour relationship when a connection is established and hello PDUs exchanged. A malfunctioning IS might, however, report another IS to be a neighbour when in fact it is not. To detect this class of failure the decision process checks that each link reported as "up" in a LSP is so reported by both Intermediate systems. If an Intermediate system considers a link "down" it shall **not** mention the link in its Link State PDUs.

On broadcast subnetworks, this class of failure shall be detected by the designated IS, which has the responsibility to ascertain the set of Intermediate systems that can all communicate on the subnetwork. The designated IS shall include these Intermediate systems (and no others) in the Link State PDU it generates for the pseudonode representing the broadcast subnetwork.

7.2.5 Multiple LSPs for the same system

The Update process is capable of dividing a single logical LSP into a number of separate PDUs for the purpose of conserving link bandwidth and processing (see 7.3.4). The Decision Process, on the other hand, shall regard the LSP with LSP Number zero in a special way. If the LSP with LSP Number zero and remaining lifetime > 0 is not present for a particular system then the Decision Process shall not process any LSPs with non-zero LSP Number which may be stored for that system.

The following information shall be taken only from the LSP with LSP Number zero. Any values which may be present in other LSPs for that system shall be disregarded by the Decision Process.

- a) The setting of the LSP Database Overload bit.
- b) The value of the IS Type field.
- c) The Area Addresses option field.
- d) the setting of the AttachedFlag bit

7.2.6 Routing algorithm overview

7.2.6.1 The routing algorithm used by the Decision Process is a *shortest path first (SPF)* algorithm. Instances of the algorithm are run independently and concurrently by all Intermediate systems in a routing domain. Intra-Domain routing of a PDU occurs on a hop-by-hop basis: that is, the algorithm determines only the next hop, not the complete path, that a data PDU will take to reach its destination. To guarantee correct and consistent route computation by every Intermediate system in a routing domain, this International Standard depends on the following properties:

- a) All Intermediate systems in the routing domain converge using identical topology information; and
- b) Each Intermediate system in the routing domain generates the same set of routes from the same input topology and set of metrics.

The first property is necessary in order to prevent inconsistent, potentially looping paths. The second property is necessary to meet the goal of determinism stated in 6.6.

7.2.6.2 A system executes the SPF algorithm to find a set of legal paths to a destination system in the routing domain. The set may consist of

- a) a single path of minimum metric sum: these are termed *minimum cost paths*;
- b) a set of paths of equal minimum metric sum: these are termed *equal minimum cost paths*; or
- c) a set of paths which will get a PDU closer to its destination than the local system: these are called *downstream paths*.

Paths which do not meet the above conditions are illegal and shall not be used. Paths whose metric sum exceeds the value of the architectural constant **MaxPathMetric** (see table 2) are also illegal and shall not be used.

7.2.6.3 The Decision Process, in determining its paths, also ascertains the identity of the adjacency which lies on the first hop to the destination on each path. These adjacencies are used to form the Forwarding Database, which the forwarding process uses for relaying PDUs.

7.2.6.4 Separate route calculations are made for each pairing of a level in the routeing hierarchy (i.e. L1 and L2) with a supported routeing metric. Since there are four routeing metrics and two levels some systems may execute multiple instances of the SPF algorithm. For example,

- if an IS is a L2 Intermediate system which supports all four metrics and computes minimum cost paths for all metrics, it would execute the SPF calculation eight times.
- if an IS is a L1 Intermediate system which supports all four metrics, and additionally computes downstream paths, it would execute the algorithm $4 \times (\text{number of neighbours} + 1)$ times.

Any implementation of an SPF algorithm meeting both the static and dynamic conformance requirements of clause 12 of this International Standard may be used. Recommended implementations are described in detail in annex C.

7.2.7 Removal of excess paths

When there are more than **maximumPathSplits** legal paths to a destination, this set shall be pruned until only **maximumPathSplits** remain. The Intermediate system shall discriminate based upon:

NOTE 15 The precise precedence among the paths is specified in order to meet the goal of determinism defined in 6.6.

- **adjacency type:** Paths associated with End system or level reachable address prefix adjacencies are retained in preference to other adjacencies
- **metric sum:** Paths having a lesser metric sum are retained in preference to paths having a greater metric sum. By metric sum it is understood, the sum of the metrics along the path to the destination.
- **neighbour ID:** where two or more paths are associated with adjacencies of the same type, an adjacency with a lower neighbour ID is retained in preference to an adjacency with a higher neighbour id.
- **circuit ID:** where two or more paths are associated with adjacencies of the same type, and same neighbour ID, an adjacency with a lower *circuit ID* is retained in preference to an adjacency with a higher *circuit ID*, where *circuit ID* is the value of
 - **ptPtCircuitID** for non-broadcast circuits,
 - **l1CircuitID** for broadcast circuits when running the Level 1 Decision Process, and
 - **l2CircuitID** for broadcast circuits when running the Level 2 Decision Process.
- **IANAAddress:** where two or more adjacencies are of the same type, same neighbour ID, and same circuit ID (e.g. where a neighbouring system has multiple LAN adapters on the same LAN) an adjacency with a lower **IANAAddress** is retained in preference to an adjacency with a higher **IANAAddress**.

7.2.8 Robustness checks

7.2.8.1 Computing routes through overloaded Intermediate systems

The Decision Process shall not utilise a link to an Intermediate system neighbour from an IS whose LSPs have the LSP Database Overload indication set. Such paths may introduce loops since the overloaded IS does not have a complete routeing information base. The Decision Process shall, however utilise the link to reach End system neighbours since these paths are guaranteed to be non-looping.

7.2.8.2 Two-way connectivity check

The Decision Process shall not utilise a link between two Intermediate Systems unless both ISs report the link.

NOTE 16 The check is not applicable to links to an End System.

Reporting the link indicates that it has a defined value for at least the default routeing metric. It is permissible for two end-points to report different defined values of the same metric for the same link. In this case, routes may be asymmetric.

7.2.9 Construction of a forwarding database

The information that is needed in the forwarding database for routeing metric k is the set of adjacencies for each system N .

7.2.9.1 Identification of nearest level 2 IS by a level 1 IS

Level 1 Intermediate systems need one additional piece of information per routeing metric: the next hop to the nearest level 2 Intermediate system according to that routeing metric. A level 1 IS shall ascertain the set, R , of “attached” level 2 Intermediate system(s) for metric k such that the total cost to R for metric k is minimal.

If there are more adjacencies in this set than `maximumPathSplits`, then the IS shall remove excess adjacencies as described in 7.2.7.

7.2.9.2 Setting the attached flag in level 2 intermediate systems

When executing the level 2 decision process for each supported metric, level 2 IS shall ascertain whether or not it can reach any destinations outside its area using that metric. The IS considers itself attached if either:

- a) it can reach at least one other area using the corresponding routeing metric, or
- b) it has at least one enabled reachable address prefix with the corresponding metric defined.

Otherwise the IS considers itself not attached.

If the IS discovers that it is not attached and `attachedFlag` was previously True, it shall:

- set `attachedFlag` for metric k to False;
- regenerate its Level 1 LSP with LSP number zero; and
- compute the “nearest level 2 Intermediate system” for metric k for insertion in the appropriate forwarding database, according to the algorithm described in 7.2.9.1 for level 1 Intermediate systems.

NOTE 17 `attachedFlag` for each metric k is examined by the Update Process, so that it will report the value in the ATT field of its Link State PDUs.

If the IS discovers that it is attached and `attachedFlag` was previously False, it shall:

- set `attachedFlag` for metric k to “True”;
- regenerate its Level 1 LSP with LSP number zero; and
- set the level 1 forwarding database entry for metric k which corresponds to “nearest level 2 Intermediate system” to “Self”.

7.2.10 Information for repairing partitioned areas

An area may become partitioned as a result of failure of one or more links in the area. However, if each of the partitions has a connection to the level 2 subdomain, it is possible to repair the partition via the level 2 subdomain, provided that the level 2 subdomain itself is not partitioned. This is illustrated in figure 5.

All the systems A–I, R and P are in the same area n . When the link between D and E is broken, the area becomes partitioned. Within each of the partitions the *Partition Designated Level 2 Intermediate system* is selected from among the level 2 Intermediate systems in that partition. In the case of partition 1 this is P, and in the case of partition 2 this is R. The level 1 repair path is then established between these two level 2 Intermediate systems. Note that the repaired link is now between P and R, not between D and E.




Figure 5 - Repair of partitioned level 1 area

The Partition Designated Level 2 Intermediate Systems repair the partition by forwarding NPDUs destined for other partitions of the area through the level 2 subdomain. They do this by acting in their capacity as Level 1 Intermediate Systems and advertising in their Level 1 LSPs adjacencies to each Partition Designated Level 2 Intermediate System in the area. This adjacency is known as a “Virtual Adjacency” or “Virtual Link”. Thus other Level 1 Intermediate Systems in a partition calculate paths to the other partitions through the Partition Designated Level 2 Intermediate System. A Partition Designated Level 2 Intermediate System forwards the Level 1 NPDUs through the level 2 subdomain by encapsulating them in ISO 8473 Data NPDUs with its Virtual Network Entity Title as the source NSAP and the “adjacent” Partition Designated Level 2 Intermediate System’s Virtual Network Entity Title as the destination NSAP. The following sub-clauses describe this in more detail.

7.2.10.1 Partition detection and virtual level 1 link creation

Partitions of a Level 1 area are detected by the Level 2 Intermediate System(s) operating within the area. Participation in the partition repair process by a Level 2 Intermediate system is predicated on the fact that all L2 ISs also function as L1 ISs within their own area. A partition of a given area exists whenever two or more Level 2 ISs located in that area are reported in the L2 LSPs as being a “Partition Designated Level 2 IS”. Conversely, when only one Level 2 IS in an area is reported as being the “Partition Designated Level 2 IS”, then that area is not partitioned. Partition repair is accomplished by the Partition Designated Level 2 IS. The election of the Partition Designated Level 2 IS as described in the next subsection must be done before the detection and repair process can begin.

In order to repair a partition of a Level 1 area, the Partition designated Level 2 IS creates a “Virtual Network Entity” to represent the partition. The Network Entity Title for this virtual network entity shall be constructed from the first listed area address from its Level 2 Link State PDU, and the ID of the Partition Designated Level 2 IS, with a selector value of `IntradomainRoutingSelector`. The IS shall also construct a virtual link (represented by a new Virtual Adjacency managed object) to each Partition Designated Level 2 IS in the area, with the NET of the partition recorded in the Identifier attribute. The virtual links are the repair paths for the partition. They are reported by the Partition Designated Level 2 IS into the entire Level 1 area by adding the ID of each “adjacent” Partition Designated Level 2 IS to the `Intermediate System Neighbours` field of its Level 1 Link State PDU. The `Virtual Flag` shall be set “True” for these Intermediate System neighbours. The metric value for this virtual link shall be the default metric value $d(N)$ obtained from this system’s Level 2 PATHS database, where N is the “adjacent” Partition Designated Level 2 IS via the Level 2 subdomain. If the computed metric value exceeds the value of `MaxLinkMetric`, the IS shall report the value of `MaxLinkMetric` for the virtual link.

An Intermediate System which operates as the Partition Designated Level 2 Intermediate System shall perform the following steps after completing the Level 2 shortest path computation in order to detect partitions in the Level 1 area and create repair paths:

- Examine Level 2 Link State PDUs of all Level 2 Inter-mediate systems. Search `Area Addresses` for any address that matches any of the addresses in `partitionAreaAddresses`. If a match is found, and the Partition Designated Level 2 Intermediate system’s ID does not equal this system’s ID, then inform the level 1 update process at this system of the

identity of the Partition Designated Level 2 Intermediate system, together with the path cost for the default routing metric to that Intermediate system.

- b) Continue examining Level 2 LSPs until all Partition Designated Level 2 Intermediate systems in other partitions of this area are found, and inform the Level 1 Update Process of all of the other Partition Designated Level 2 Intermediate systems in other partitions of this area, so that
 - 1) Level 1 Link State PDUs can be propagated to all other Partition designated level 2 Intermediate systems for this area (via the level 2 subdomain).
 - 2) All the Partition Designated Level 2 Intermediate systems for other partitions of this area can be reported as adjacencies in this system's Level 1 Link State PDUs.

If a partition has healed, the IS shall destroy the associated virtual network entity and virtual link by deleting the Virtual Adjacency. The Partition Designated Level 2 IS detects a healed partition when another Partition Designated Level 2 IS listed as a virtual link in its Level 1 Link State PDU was not found after running the partition detection and virtual link creation algorithm described above.

If such a Virtual Adjacency is created or destroyed, the IS shall generate a `partitionVirtualLinkChange` event.

7.2.10.2 Election of partition designated level 2 intermediate system

The Partition Designated Level 2 IS is a Level 2 IS which:

- reports itself as "attached" by the default metric in its LSPs;
- reports itself as implementing the partition repair option;
- operates as a Level 1 IS in the area;
- is reachable from the IS performing the calculation via Level 1 routing without traversing any virtual links; and
- has the lowest ID

The election of the Partition Designated Level 2 IS is performed by running the decision process algorithm after the Level 1 decision process has finished, and before the Level 2 decision process to determine Level 2 paths is executed.

In order to guarantee that the correct Partition Designated Level 2 IS is elected, the decision process is run using only the Level 1 LSPs for the area, and by examining only the **Intermediate System Neighbours** whose **Virtual Flag** is FALSE. The results of this decision process is a set of all the Level 1 Intermediate Systems in the area that can be reached via Level 1, non-virtual link routing. From this set, the Partition Designated Level 2 IS is selected by choosing the IS for which:

- **IS Type** (as reported in the Level 1 LSP) is "Level 2 Intermediate System";
- **ATT** indicates "attached" by the default metric;
- **P** indicates support for the partition repair option; and
- **ID** is the lowest among the subset of attached Level 2 Intermediate Systems.

7.2.10.3 Computation of partition area addresses

For systems which do not implement partition repair, the value of `partitionAreaAddresses` is identical to the value computed for `areaAddresses` as described in 7.2.11. For systems which do implement partition repair, the value of `partitionAreaAddresses` is determined as follows. A Level 2 Intermediate System shall compute the set of `partitionAreaAddresses`. The set of `partitionAreaAddresses` is defined as the union of all `manualAreaAddresses` as reported in the Level 1 link state PDUs of all level 1 and level 2 Intermediate Systems reachable in the partition by the traversal of non-virtual links. If more than `maximumAreaAddresses` are present, the Intermediate system shall retain only those areas with numerically lowest area address (as described in 7.1.8). If one of the local system's `manualAreaAddresses` is so rejected the event `manual-AddressDroppedFromArea` shall be generated.

7.2.10.4 Encapsulation of NPDU's across the virtual link

All NPDU's sent over virtual links shall be encapsulated as ISO 8473 Data NPDU's. The IS shall encapsulate NPDU's in ISO 8473 Data PDU's as follows:

- a) The **Source Address** field of the encapsulating Data NPDU shall contain the Virtual Network Entity Title of the Partition Designated Level 2 IS that is forwarding the NPDU over the virtual link
- b) The **Destination Address** field shall contain the Virtual NET of the "adjacent" Partition Designated Level 2 IS

- c) The **SEL** field in both addresses shall contain the **IntradomainRouteingSelector** value.
- d) If the PDU to be encapsulated is an ISO 8473 Data or Error PDU the IS shall,
 - 1) Ensure that the header of the encapsulating PDU indicates routeing by the default metric. Routeing by the default metric may be indicated either by including a QoS parameter with its bits set in accordance with Table 1, or by carrying no QoS parameter at all [see 7.4.2 item a)]
 - 2) copy the **Segmentation Permitted**, and **Error Report** flags from the inner PDU.
 - 3) copy the **Lifetime** field from the inner NPDU. When the inner NPDU is decapsulated, its **Lifetime** field shall be set to the value of the **Lifetime** field in the outer NPDU.
- e) If the PDU to be encapsulated is an ISO 8473 Data PDU, the IS shall not segment it before encapsulation, unless the total length of the Data NPDU (including header) exceeds 65512 octets. In that case, the original Data NPDU shall first be segmented, then encapsulated. In all cases, the encapsulated Data NPDU may need to be segmented by ISO 8473 before transmission in which case it must be reassembled and decapsulated by the destination Partition Designated Level 2 IS.
- f) If the PDU to be encapsulated is an LSP or SNP, the IS shall:
 - 1) set the **Segmentation Permitted** flag to "True";
 - 2) set the **Error Report** flag to "False"
 - 3) set the **Lifetime** field to 255. When an inner LSP is decapsulated, its remaining lifetime shall be decremented by half the difference between 255 and the value of the **Lifetime** field in the outer NPDU.

The encapsulation is further described as part of the forwarding process in 7.4.3.2. The decapsulation is described as part of the Receive process in 7.4.4.

7.2.10.5 Encapsulated NPDU Forwarding Mechanism

A level 2 Intermediate System (including a level 2 IS which does not support the partition repair function) over a virtual link has to forward the encapsulated NPDUs to a destination of partition Designated Level 2 IS. To achieve this, the level 2 IS shall perform the following functions:

- a) It shall construct the level 2 forwarding database consisting of the pairs <NET of partition Designated level 2 IS, adjacency to the partition Designated level 2 IS>. The NET of the partition Designated level 2 Intermediate System is constructed from first listed area address from the **Area Addresses** field and the **Partition Designated Level 2 Intermediate System** field reported by a level 2 LSP stored in a level 2 Link State database, with a selector value of **Intradomain-RouteingSelector**.
- b) It shall choose the level 2 route with the longest address prefix match.

7.2.11 Computation of area addresses

A Level 1 or Level 2 Intermediate System shall compute the values of **areaAddresses** (the set of area addresses for this Level 1 area), by forming the union of the sets of **manualAreaAddresses** reported in the **Area Addresses** field of all Level 1 LSPs with LSP number zero in the local Intermediate system's link state database from all ISs which are reachable via Level 1 routeing.

NOTE 18 This includes all source systems, whether currently reachable via Level 1 links - both virtual and non-virtual. It also includes the local Intermediate System's own Level 1 LSP with LSP number zero.

NOTE 19 There is no requirement for this set to be updated immediately on each change to the database contents. It is permitted to defer the computation until the next running of the Decision Process.

If more than **maximumAreaAddresses** are present, the Intermediate system shall retain only those areas with numerically lowest area address (as described in 7.1.8). If one of the local system's **manualAreaAddresses** is rejected the event **manualAddressDroppedFromArea** shall be generated.

7.2.12 Order of preference of routes

7.2.12.1 If an Intermediate system takes part in level 1 routeing, and determines (by looking at the area address) that a given destination is reachable within its area, then that destination will be reached exclusively by use of level 1 routeing. In particular:

- a) Level 1 routing is always based on internal metrics.
- b) Amongst routes in the area, routes on which the requested QoS (if any) is supported are always preferred to routes on which the requested QoS is not supported.
- c) Amongst routes in the area of the same QoS, the shortest routes are preferred. For determination of the shortest path, if a route with specific QoS support is available, then the specified QoS metric is used, otherwise the default metric is used.

7.2.12.2 If an Intermediate system takes part in level 1 routing, does not take part in level 2 routing, and determines (by looking at the area address) that a given destination is *not* reachable within its area, and at least one attached level 2 IS is reachable in the area, then that destination will be reached by routing to a level 2 Intermediate system as follows:

- a) Level 1 routing is always based on internal metrics.
- b) Amongst routes in the area to attached level 2 ISs, routes on which the requested QoS (if any) is supported are always preferred to routes on which the requested QoS is not supported.
- c) Amongst routes in the area of the same QoS to attached level 2 ISs, the shortest route is preferred. For determination of the shortest path, if a route on which the specified QoS is available, then the specified QoS metric is used, otherwise the default metric is used.

7.2.12.3 If an Intermediate system takes part in level 2 routing and the IS determines (by looking at the area address) that the destination is reachable within its area, then the destination will be reached either by level 1 or level 2 routing, as follows:

- a) level 1 routing is always based on internal metrics;
- b) a level 2 route, if it exists, shall be preferred over a level 1 route that traverses a virtual link. Otherwise, the level 1 route shall be preferred;
- c) amongst routes in the area, routes on which the requested QoS (if any) is supported are always preferred to routes on which the requested QoS is not supported;
- d) amongst routes in the area of the same QoS, the shortest routes are preferred. For determination of the shortest path, if a route with specific QoS support is available, then the specified QoS metric is used, otherwise the default metric is used.

7.2.12.4 If an Intermediate system takes part in level 2 routing and is attached, and the IS determines (by looking at the area address) that a given destination is not reachable within its area, forwarding for that destination will select routes as follows:

- a) Routes on which the requested QoS (if any) is supported are always preferred to routes on which the requested QoS is not supported.
- b) Amongst routes of the same QoS, routes are prioritised as follows:
 - 1) Highest precedence: routes whose destination area address matches the area address of an area inside the routing domain (i.e. the route does not go outside the routing domain). In case of multiple matching address prefixes the longest prefix shall be preferred.
 - 2) Medium precedence: routes constructed from the Reachable Address Prefix information in an LSP which indicates an internal metric. In the case of multiple prefixes which match a given destination address which all have internal metrics, then the longest prefix shall be preferred.
 - 3) Lowest precedence: routes constructed from the Reachable Address Prefix information in an LSP which indicates an external metric. In the case of multiple prefixes which match a given destination address which all have external metrics, then the longest prefix shall be preferred.
- c) For routes with equal precedence as specified above, the shortest path shall be preferred. For determination of the shortest path, a route supporting the specified QoS is used if available; otherwise a route using the default metric shall be used. Amongst routes of equal cost, load splitting may be performed.

7.3 Update process

The Update process is responsible for generating and propagating Link State information reliably throughout the routing domain.

The Link State information is used by the Decision process to calculate routes.

7.3.1 Input and output

INPUT

- Adjacency Database – maintained by the Subnetwork Dependent Functions
- Reachable Address managed objects - maintained by System Management
- Notification of Adjacency Database Change – notification by the Subnetwork Dependent Functions that an adjacency has come up, gone down, or changed cost. (Circuit up, Circuit down, Adjacency Up, Adjacency Down, and Cost change events)
- **AttachedFlag** – (level 2 Intermediate systems only), a flag computed by the Level 2 Decision Process indicating whether this system can reach (via level 2 routeing) other areas
- Link State PDUs – The Receive Process passes Link State PDUs to the Update Process, along with an indication of which adjacency it was received on.
- Sequence Numbers PDUs – The Receive Process passes Sequence Numbers PDUs to the Update Process, along with an indication of which adjacency it was received on.
- Other Partitions – The Level 2 Decision Process makes available (to the Level 1 Update Process on a Level 2 Intermediate system) a list of Partition Designated Level 2 Intermediate system, Level 2 default metric value pairs, for other partitions of this area.

OUTPUT

- Link State Database
- Signal to the Decision Process of an event, which is either the receipt of a Link State PDU with different information from the stored one, or the purging of a Link State PDU from the database. The reception of a Link State PDU which has a different sequence number or Remaining Lifetime from one already stored in the data base, but has an identical variable length portion, shall not cause such an event.

NOTE 20 An implementation may compare the checksum of the stored Link State PDU, modified according to the change in sequence number, with the checksum of the received Link State PDU. If they differ, the implementation may assume that the variable length portions are different and it shall signal an event to the Decision Process. However, if the checksums are the same, an octet for octet comparison is needed determine whether or not to signal the event.

7.3.2 Generation of local link state information

The Update Process is responsible for constructing a set of Link State PDUs. The purpose of these Link State PDUs is to inform all the other Intermediate systems (in the area, in the case of Level 1, or in the Level 2 subdomain, in the case of Level 2), of the state of the links between the Intermediate system that generated the PDUs and its neighbours.

The Update Process in an Intermediate system shall generate one or more new Link State PDUs under the following circumstances:

- a) upon timer expiration;
- b) when notified by the Subnetwork Dependent Functions of an Adjacency Database Change;
- c) when a change to some System Management parameter would cause the information in the LSP to change (for example, a change in `manualAreaAddresses`).

7.3.3 Use of manual routeing information

Manual routeing information is routeing information entered by System Management. It may be specified in two forms.

- a) *Manual Adjacencies*
- b) *Reachable Addresses*

These are described in the following sub-clauses.

7.3.3.1 Manual adjacencies

An End system adjacency may be created by System Management. Such an adjacency is termed a *manual End system adjacency*. In order to create a manual End system adjacency, system management shall specify:

- a) the (set of) system IDs reachable over that adjacency; and
- b) the corresponding SNPA Address.

A manual adjacency is considered to be “active” when all the following conditions are true:

- a) The **operationalState** of the linkage managed object associated with the parent circuit is “Enabled”
- b) the **adjacencyState** of the managed object associated with the adjacency is “On” and
- c) the parent circuit is of type broadcast or its underlying data link is operational.

Whenever a manual adjacency changes from being “inactive” to “active” a signal shall be generated to the Update process to cause it to include the IDs of the manual adjacency in the Level 1 LSPs generated by that system as described in 7.3.7.

Whenever a manual adjacency changes from being “active” to “inactive”, a signal shall be generated to the Update process to cause it to cease including the IDs of the manual adjacency in the Level 1 LSPs.

NOTE 21 Manual End system adjacencies are not included in a Level 1 LSPs issued on behalf of a pseudonode, since that would presuppose that all Intermediate systems on a broadcast subnetwork had the same set of manual adjacencies as defined for this circuit.

Metrics assigned to Manual adjacencies shall be internal metrics.

7.3.3.2 Reachable addresses

A Level 2 Intermediate system may have a number of **reachableAddress** managed objects created by System management. When a **reachableAddress** managed object is in **operationalState** “Enabled” and the **linkage** managed object associated with its parent Circuit is also in **operationalState** “Enabled”, the name and each of its defined routing metrics shall be included in Level 2 LSPs generated by this system.

Metrics assigned to **reachableAddress** managed objects may be either Internal or External.

A reachable address is considered to be “active” when all the following conditions are true:

- a) The **operationalState** of the **linkage** managed object associated with the parent circuit is “Enabled”;
- b) the **operationalState** of the **reachableAddress** is “Enabled”; and
- c) the parent circuit is of type broadcast or its underlying data link is operational.

Whenever a reachable address changes from being “inactive” to “active” a signal shall be generated to the Update process to cause it to include the Address Prefix of the reachable address and each of its defined routing metrics in the Level 2 LSPs generated by that system as described in 7.3.9.

Whenever a reachable address changes from being “active” to “inactive”, a signal shall be generated to the Update process to cause it to cease including the Address Prefix of the reachable address in the Level 2 LSPs.

7.3.4 Multiple LSPs

7.3.4.1 Because a Link State PDU is limited in size to **ReceiveLSPBufferSize**, it may not be possible to include information about all of a system’s neighbours in a single LSP. In such cases, a system may use multiple LSPs to convey this information. Each LSP in the set carries the same **sourceID** field (see clause 9), but sets its own **LSP Number** field individually. Each of the several LSPs is handled independently by the Update Process, thus allowing distribution of topology updates to be pipelined. However, the Decision Process recognises that they all pertain to a common originating system because they all use the same **sourceID**.

NOTE 22 Even if the amount of information is small enough to fit in a single LSP, a system may optionally choose to use several LSPs to convey it; use of a single LSP in this situation is not mandatory.

NOTE 23 In order to minimize the transmission of redundant information, it is advisable for an IS to group Reachable Address Prefix information by the circuit with which it is associated. Doing so will ensure that the minimum number of LSP fragments need be transmitted if a circuit to another routing domain changes state.

7.3.4.2 The maximum sized Level 1 or Level 2 LSP which may be generated by a system is controlled by the values of the management parameters `originatingL1LSPBufferSize` or `originatingL2LSPBufferSize` respectively.

NOTE 24 These parameters should be set consistently by system management. If this is not done, some adjacencies will fail to initialise.

7.3.4.3 The IS shall treat the LSP with LSP Number zero in a special way, as follows:

- a) The following fields are meaningful to the decision process only when they are present in the LSP with LSP Number zero:
 - 1) The setting of the **LSP Database Overload** bit.
 - 2) The value of the **IS Type** field.
 - 3) The **Area Addresses** option field. The **Area Addresses** option shall be present only in an LSP with LSP number zero. If there are more area addresses than will fit in a single instance of the **Area Addresses** option field, then the IS shall place 12 area addresses in each instance of the field except the last.
 - 4) the setting of the **AttachedFlag** bit
- b) When the values of any of the above items are changed, an Intermediate System shall reissue the LSP with LSP Number zero, to inform other Intermediate Systems of the change. Other LSPs need not be reissued.

7.3.4.4 Once a particular adjacency has been assigned to a particular LSP Number, it is desirable that it not be moved to another LSP Number. This is because moving an adjacency from one LSP to another can cause temporary loss of connectivity to that system. This can occur if the new version of the LSP which originally contained information about the adjacency (which now does not contain that information) is propagated before the new version of the other LSP (which now contains the information about the adjacency).

NOTE 25 An implementation is recommended to ensure that the number of LSPs generated for a particular system is within approximately 10 % of the optimal number which would be required if all LSPs were densely packed with neighbour options. Where possible this should be accomplished by re-using space in LSPs with a lower LSP Number for new adjacencies.

If it is necessary to move an adjacency from one LSP to another, the **SRMflags** (see 7.3.15) for the two new LSPs shall be set as an atomic action.¹⁾

7.3.4.5 When some event requires changing the LSP information for a system, the system shall reissue that (or those) LSPs which would have different contents. It is not required to reissue the unchanged LSPs. Thus a single End system adjacency change only requires the reissuing of the LSP containing the **End System Neighbours** option referring to that adjacency.

7.3.4.6 If an LSP becomes empty, because all of the adjacencies reported in that LSP no longer exist, it is recommended that an IS purge that LSP. An IS may alternatively choose to reissue it as an empty LSP (i.e. with no options), although purging is preferable. An IS should not allow the LSP to age out as this results in having incorrect routing information in all of the LSP databases.

7.3.5 Periodic LSP generation

The Update Process shall periodically re-generate and propagate on every circuit with an IS adjacency of the appropriate level (by setting **SRMflag** on each circuit), all the LSPs (Level 1 and/or Level 2) for the local system and any pseudonodes for which it is responsible. The Intermediate system shall regenerate each LSP at intervals of at most `maximumLSPGenerationInterval`, with jitter applied as described in 10.1.

It is not required to synchronise the regeneration of the individual LSPs.

¹⁾ If the two **SRMflags** are not set atomically, a race condition will exist in which one of the two LSPs may be propagated quickly, while the other waits for an entire propagation cycle. If this occurs, adjacencies will be falsely eliminated from the topology and routes may become unstable for periods of time potentially as large as `maximumLSPGenerationInterval`.

7.3.6 Event driven LSP Generation

In addition to the periodic generation of LSPs, an Intermediate system shall generate an LSP when an event occurs which would cause the information content to change. The following events may cause such a change.

- an Adjacency or Circuit Up/Down event
- a change in Circuit metric
- a change in Reachable Address metric
- a change in `manualAreaAddresses`
- a change in `systemID`
- a change in Designated Intermediate System status
- a change in the waiting status
- creation or delete of a virtual circuit
- a changed of `AttachedFlag`

When such an event occurs the IS shall re-generate changed LSP(s) with a new sequence number. If the event necessitated the generation of an LSP which had not previously been generated (for example, an adjacency “Up” event for an adjacency which could not be accommodated in an existing LSP), the sequence number shall be set to one. The IS shall then propagate the LSP(s) on every circuit by setting `SRMflag` for each circuit. The timer `maximumLSPGenerationInterval` shall **not** be reset.

There is a hold down timer (`minimumLSPGenerationInterval`) on the generation of each individual LSP.

7.3.7 Generation of level 1 LSPs (non-pseudonode)

The Level 1 Link State PDU not generated on behalf of a pseudonode contains the following information in its variable length fields.

- In the `Area Addresses` option — the set of `manualAreaAddresses` for this Intermediate System.
- In the `Intermediate System Neighbours` option — the set of Intermediate system IDs of neighbouring Intermediate systems formed from
- The set of `neighbourSystemIDs` with an appended zero octet (indicating non-pseudonode) from adjacencies in the state “Up”, on circuits of type “ptToPt”, “staticIn”, or “staticOut”, with
 - x `neighbourSystemType` “L1 Intermediate System”
 - x `neighbourSystemType` “L2 Intermediate System” and `adjacencyUsage` “Level 1” or “Level 1 and 2”.

The metrics shall be set to the values of Level 1 metric k of the circuit for each supported routing metric.

- The set of `l1CircuitIDs` for all circuits of type “Broadcast” (i.e. the neighbouring pseudonode IDs) .
The metrics shall be set to the values of Level 1 metric _{k} of the circuit for each supported routing metric, where k = one of (default, delay, error, expense).
- The set of IDs with an appended zero octet derived from the Network Entity Titles of all Virtual Adjacencies of this IS. (Note that the `Virtual Flag` is set when encoding these entries in the LSP — see 7.2.10.) The default metric shall be set to the total cost to the virtual NET for the default routing metric. The remaining metrics shall be set to the value indicating *unsupported*.
- In the `End System Neighbours` option — the set of Ids of neighbouring End systems formed from
- The `systemID` of the Intermediate System itself, with a value of zero for all supported metrics.
- The set of `neighbourSystemIDs` from all automatically created adjacencies, in state “Up”, on circuits of type “ptToPt”, “staticIn”, or “staticOut”, with `neighbourSystemType` “End system”. The metrics shall be set to the values of Level 1 metric _{k} of the circuit for each supported routing metric, where k = one of (default, delay, error, expense).
- The set of `neighbourSystemIDs` from all system management-created adjacencies in state “Up”, on all circuits, with `neighbourSystemType` “End system” The metrics shall be set to the values of Level 1 metric _{k} of the circuit for each supported routing metric, where k = one of (default, delay, error, expense).

- In the Authentication Information field — if the system's `areaTransmitPassword` is non-null, include the Authentication Information field containing an Authentication Type of "Password", and the value of the `areaTransmitPassword`.

7.3.8 Generation of level 1 pseudonode LSPs

An IS shall generate a Level 1 pseudonode Link State PDU for each circuit for which this Intermediate System is the Level 1 LAN Designated Intermediate System. The LSP shall specify the following information in its variable length fields. In all cases a value of zero shall be used for all supported routing metrics

- The Area Addresses option is not present.

NOTE 26 This information is not required since the set of area addresses for the node issuing the pseudonode LSP will already have been made available via its own non-pseudonode LSP.

- In the Intermediate System Neighbours option — the set of Intermediate System IDs of neighbouring Intermediate Systems on the circuit for which this pseudonode LSP is being generated formed from
 - The Designated Intermediate System's own `systemID` with an appended zero octet (indicating non-pseudonode).
 - The set of `neighbourSystemIDs` with an appended zero octet (indicating non-pseudonode) from adjacencies on this circuit in the `adjacencyState` "Up", with
 - x `neighbourSystemType` "L1 Intermediate System"
 - x `neighbourSystemType` "L2 Intermediate System" and `adjacencyUsage` "Level 1".
- In the End System Neighbours option — the set of IDs of neighbouring End systems formed from:
 - The set of `neighbourSystemIDs` from all automatically-created adjacencies, in state "Up", on the circuit for which this pseudonode is being generated, with `neighbourSystemType` "End-system".
- In the Authentication Information field — if the system's `areaTransmitPassword` is non-null, include the Authentication Information field containing an Authentication Type of "Password", and the value of the `areaTransmitPassword`.

7.3.9 Generation of level 2 LSPs (non-pseudonode)

The Level 2 Link State PDU not generated on behalf of a pseudonode contains the following information in its variable length fields:

- In the Area Addresses option — the set of `partitionAreaAddresses` for this intermediate system as described in 7.2.10.3
- In the Partition Designated Level 2-IS option — the ID of the Partition Designated Level 2 Intermediate System for the partition.
- In the Intermediate System Neighbours option — the set of Intermediate system IDs of neighbouring Intermediate systems formed from:
 - The set of `neighbourSystemIDs` with an appended zero octet (indicating non-pseudonode) from adjacencies in the state "Up", on circuits of type "ptToPt", "staticIn", or "staticOut", with `neighbourSystemType` "L2 Intermediate System".
 - The set of `I2CircuitIDs` for all circuits whose linkage managed object is of type "Broadcast". (i.e. the neighbouring pseudonode IDs)

The metric and metric type shall be set to the values of Level 2 metric_k of the circuit for each supported routing metric.

- In the Prefix Neighbours option — the set of variable length prefixes formed from
 - The set of `prefixAddresses` from all `reachableAddress` managed objects in `operationalState` "Enabled", on all circuits in state "Up".

The metrics shall be set to the values of Level 2 metric_k for the reachable address.

- In the Authentication Information field — if the system's `domainTransmitPassword` is non-null, include the Authentication Information field containing an Authentication Type of "Password", and the value of the `domainTransmitPassword`.

7.3.10 Generation of level 2 pseudonode LSPs

A Level 2 pseudonode Link State PDU is generated for each circuit for which this Intermediate System is the Level 2 LAN Designated Intermediate System and contains the following information in its variable length fields. In all cases a value of zero shall be used for all supported routing metrics.

- The **Area Addresses** option is not present.

NOTE 27 This information is not required since the set of area addresses for the node issuing the pseudonode LSP will already have been made available via its own non-pseudonode LSP.

- In the **Intermediate System Neighbours** option — the set of Intermediate System IDs of neighbouring Intermediate Systems on the circuit for which this pseudonode LSP is being generated formed from:
 - The Designated Intermediate System's own **systemID** with an appended zero octet (indicating non-pseudonode).
 - The set of **neighbourSystemIDs** with an appended zero octet (indicating non-pseudonode) from adjacencies on this circuit in the state "Up" with **neighbourSystemType** "L2 Intermediate System".
- The **Prefix Neighbours** option is not present.
- In the **Authentication Information** field — if the system's **domainTransmitPassword** is non-null, include the **Authentication Information** field containing an **Authentication Type** of "Password", and the value of the **domainTransmitPassword**.

7.3.11 Generation of the checksum

This International Standard makes use of the checksum function defined in ISO 8473.

The source IS shall compute the LSP Checksum when the LSP is generated. The checksum shall never be modified by any other system. The checksum allows the detection of memory corruptions and thus prevents both the use of incorrect routing information and its further propagation by the Update Process.

The checksum shall be computed over all fields in the LSP which appear after the **Remaining Lifetime** field. This field (and those appearing before it) are excluded so that the LSP may be aged by systems without requiring recomputation.

As an additional precaution against hardware failure, when the source computes the Checksum, it shall start with the two checksum variables (C0 and C1) initialized to what they would be after computing for the **systemID** portion of its Source ID. (This value is computed and stored when the Network entity is enabled and whenever **systemID** changes.) The IS shall then resume Checksum computation on the contents of the PDU after the first ID Length octets of the **Source ID** field.

NOTE 28 All Checksum calculations on the LSP are performed treating the **Source ID** field as the first octet. This procedure prevents the source from accidentally sending out Link State PDUs with some other system's ID as source.

7.3.12 Initiating transmission

The IS shall store the generated Link State PDU in the Link State Database, overwriting any previous Link State PDU with the same LSP Number generated by this system. If the generated Link State PDU is of type level 1, the Intermediate System shall set the **SRMflag** for that LSP over each circuit that has at least one adjacency of **neighborSystemType** "L1 Intermediate System", or at least one adjacency of **neighborSystemType** "L2 Intermediate System" and **adjacencyUsage** "Level 1" or "Level 1 and 2", and all virtual links. If the generated Link State PDU is of type level 2, the Intermediate System shall set the **SRMflag** for that LSP over each circuit which has at least one adjacency of **neighborSystemType** "L2 Intermediate System".

An Intermediate system shall ensure (by reserving resources, or otherwise) that it will always be able to store and internalise its own non-pseudonode zero th LSP. In the event that it is not capable of storing and internalising one of its own LSPs it shall enter the overloaded state as described in 7.3.19.1.

NOTE 29 It is recommended that an Intermediate system ensures (by reserving resources, or otherwise) that it will always be able to store and internalise all its own (zero and non-zero, pseudonode and non-pseudonode) LSPs.

7.3.13 Preservation of order

When an existing Link State PDU is re-transmitted (with the same or a different sequence number), but with the same information content (i.e. the variable length part) as a result of there having been no changes in the local topology databases, the order of the information in the variable length part shall be the same as that in the previously transmitted LSP.

NOTE 30 If a sequence of changes results in the state of the database returning to some previous value, there is no requirement to preserve the ordering. It is only required when there have been no changes whatever. This allows the receiver to detect that there has been no change in the information content by performing an octet for octet comparison of the variable length part, and hence not re-run the decision process.

7.3.14 ropagation of LSPs

7.3.14.1 The update process is responsible for propagating Link State PDUs throughout the domain (or in the case of Level 1, throughout the area).

The basic mechanism is flooding, in which each Intermediate system propagates to all its neighbour Intermediate systems except that neighbour from which it received the PDU. Duplicates are detected and dropped.

7.3.14.2 Link state PDUs are received from the Receive Process. The maximum size control PDU (Link State PDU or Sequence Numbers PDU) which a system expects to receive shall be **ReceiveLSPBufferSize** octets. (i.e. the Update process must provide buffers of at least this size for the reception, storage and forwarding of received Link State PDUs and Sequence Numbers PDUs.) If a control PDU larger than this size is received, it shall be treated as if it had an invalid checksum (i.e. ignored by the Update Process and a **corruptedLSPReceived** event generated).

Upon receipt of a Link State PDU the Update Process shall perform the following functions:

- a) Level 2 Link State PDUs shall be propagated on circuits which have at least one Level 2 adjacency.
- b) Level 1 Link State PDUs shall be propagated on circuits which have at least one Level 1 adjacency or at least one Level 2 adjacency not marked "Level 2 only, and virtual links".
- c) When propagating a Level 1 Link State PDU on a broadcast subnetwork, the IS shall transmit to the multi-destination subnetwork address **AllL1IS**.
- d) When propagating a Level 2 Link State PDU on a broadcast subnetwork, the IS shall transmit to the multi-destination subnetwork address **AllL2IS**.

NOTE 31 When propagating a Link State PDU on a general topology subnetwork the Data Link Address is unambiguous (because Link State PDUs are not propagated across Dynamically Assigned circuits).

- e) An Intermediate system receiving a Link State PDU with an incorrect LSP Checksum or with an invalid PDU syntax shall
 - 1) generate a **corruptedLSPReceived** circuit event,
 - 2) discard the PDU.
- f) A Intermediate system receiving a Link State PDU which is new (as identified in 7.3.16) shall
 - 1) store the Link State PDU into Link State database, and
 - 2) mark it as needing to be propagated upon all circuits except that upon which it was received.
- g) When a Intermediate system receives a Link State PDU from source S, which it considers older than the one stored in the database for S, it shall set the **SRMflag** for S's Link State PDU associated with the circuit from which the older Link State PDU was received. This indicates that the stored Link State PDU needs to be sent on the link from which the older one was received.
- h) When a system receives a Link State PDU which is the same (not newer or older) as the one stored, the Intermediate system shall
 - 1) acknowledge it if necessary, as described in 7.3.17, and
 clear the **SRMflag** for that circuit for that Link State PDU.
- i) A Link State PDU received with a zero checksum shall be treated as if the Remaining Lifetime were zero. The age, if not zero, shall be overwritten with zero.

7.3.14.3 The Update Process scans the Link State Database for Link State PDUs with **SRMflag** set. When one is found, provided the timestamp **lastSent** indicates that it was propagated no more recently than **minimumLSPTransmissionInterval**, the IS shall

- a) transmit it on all circuits with **SRMflag** set, and
- b) update **lastSent**.

7.3.15 Manipulation of SRM and SSN flags

For each Link State PDU, and for each circuit over which routing messages are to be exchanged (i.e. not on DA circuits), there are two flags:

Send Routing Message (SRMflag) – if set, indicates that Link State PDU should be transmitted on that circuit. On broadcast circuits **SRMflag** is cleared as soon as the LSP has been transmitted, but on non-broadcast circuits **SRMflag** is only cleared on reception of a Link State PDU or Sequence Numbers PDU as described below.

SRMflag shall never be set for an LSP with sequence number zero, nor on a circuit whose **externalDomain** attribute is “True” (See 7.3.15.2).

If an LSP is larger than the **dataLinkBlockSize** of a circuit over which it is to be propagated, then no attempt shall be made to propagate this LSP over the circuit and an **ISPToolargeToPropagate** alarm shall be generated. This behavior may be implemented in either of the following ways:

- 1) The relative sizes of the LSP and the **dataLinkBlockSize** of the circuit over which it is to be propagated is checked at the time the **SRMflag** is to be set. If the LSP size is larger, the **SRMflag** is not set and an **ISPToolargeToPropagate** alarm is generated.
- 2) The relative sizes of the LSP and the **dataLinkBlockSize** of the circuit over which it is to be propagated is checked at the time an attempt to transmit the LSP is to be made. If the LSP size is larger the **SRMflag** is cleared and an **ISPToolargeToPropagate** alarm is generated.

Send Sequence Numbers (SSNflag) – if set, indicates that information about that Link State PDU should be included in a Partial Sequence Numbers PDU transmitted on that circuit with an associated linkage. When the Sequence Numbers PDU has been transmitted **SSNflag** is cleared. Note that the Partial Sequence Numbers PDU serves as an acknowledgement that a Link State PDU was received.

SSNflag shall never be set on a circuit whose **externalDomain** attribute is “True”.

7.3.15.1 Action on receipt of a link state PDU

When a Link State PDU is received on a circuit C, the IS shall perform the following functions

- a) Perform the following PDU acceptance tests:
 - 1) If the LSP was received over a circuit whose **externalDomain** attribute is “True”, the IS shall discard the PDU.
 - 2) If a level 1 Intermediate System receives a level 2 LSP, then the IS shall discard the LSP.
 - 3) If a level 2 Intermediate System receives a level 1 LSP on a circuit with **manualL2OnlyMode** set to “True”, then the IS shall discard the LSP.
 - 4) If the **ID Length** field of the PDU is not equal to the value of the IS’s **routingDomainIDLength**, the PDU shall be discarded and an **IDFieldLengthMismatch** event generated.
 - 5) If this is a level 1 LSP, and the **Maximum Area Addresses** field of the PDU is not equal to the value of the IS’s **maximumAreaAddresses** then the PDU shall be discarded and a **maximumAreaAddresses-Mismatch** event generated, unless the IS only implements a value of three for **maximumArea-Addresses**, in which case this check may be omitted.
 - 6) If the LSP is received on a broadcast circuit and the source subnetwork address of the LSP does not match the **neighbourSNPAAAddress** of an adjacency of the same level (e.g. a level 1 LSP with a level 1 adjacency) on the circuit over which the LSP was received, then the IS shall discard the LSP without generating an event.

If the LSP is received on a non-broadcast circuit and there is no adjacency of the same level (e.g. a level 1 LSP with a level 1 or level 1 & 2 adjacency) on the circuit over which the LSP was received, then the IS shall discard the LSP without generating an event.

- 7) If this is a level 1 LSP, and the value of **areaTransmitPassword** or the set of **areaReceive-Passwords** is non-null, then perform the following tests:
 - i. If the PDU does not contain the **Authentication Information** field then the PDU shall be discarded and an **authenticationFailure** event generated.
 - ii. If the PDU contains the **Authentication Information** field, but the **Authentication Type** is not equal to "Password", then:
 - a. If the IS implements the authentication procedure indicated by the **Authentication Type** whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - b. If the IS does not implement the authentication procedure indicated by the **Authentication Type** then the IS shall ignore the PDU and generate an **authenticationFailure** event."
 - iii. Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of **areaReceivePasswords**, augmented by the value of the **areaTransmitPassword**. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the above values, then the IS shall ignore the PDU and generate an **authenticationFailure** event.
 - 8) If this is a level 2 LSP, and the value of **domain-TransmitPassword** or the set of **domainReceive-Passwords** is non-null, then perform the following tests:
 - i. If the PDU does not contain the **Authentication Information** field then the PDU shall be discarded and an **authenticationFailure** event generated.
 - ii. If the PDU contains the **Authentication Information** field, but the **Authentication Type** is not equal to "Password", then:
 - a) If the IS implements the authentication procedure indicated by the **Authentication Type** whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - b) If the IS does not implement the authentication procedure indicated by the **Authentication Type** then the IS shall ignore the PDU and generate an **authenticationFailure** event."
 - iii. Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of **domainReceivePasswords**, augmented by the value of the **domainTransmitPassword**. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the above values, then the IS shall ignore the PDU and generate an **authenticationFailure** event.
 - 9) If this is a Level 1 LSP and the PDU is larger than the local **originatingL1LSPBufferSize** or this is a Level 2 LSP and the PDU is larger than the local **originatingL2LSP-BufferSize** then an **originatingLSPBufferSize-Mismatch** alarm shall be generated. Otherwise, if this is a Level 1 LSP and the **originatingLSPBufferSize** field is present and the value in the PDU does not match the local value of **originatingL1LSPBufferSize** or this is a Level 2 LSP and the **originatingLSPBufferSize** field is present and the value in the PDU does not match the local value of **originatingL2LSPBufferSize** then an **originatingLSP-BufferSizeMismatch** alarm shall be generated. The PDU shall be accepted for further processing.
- b) If the LSP has zero Remaining Lifetime, perform the actions described in 7.3.16.4.
 - c) If the source *S* of the LSP is an IS or pseudonode for which all but the last octet are equal to the **systemID** of the receiving Intermediate System, and the receiving Intermediate System does not have that LSP in its database, or has that LSP, but no longer considers it to be in the set of LSPs generated by this system (e.g. it was generated by a previous incarnation of the system), then initiate a network wide purge of that LSP as described in 7.3.16.4.
 - d) If the source *S* of the LSP is a system (pseudonode or otherwise) for which the first ID Length octets are equal to the **systemID** of the receiving Intermediate system, and the receiving Intermediate system has an LSP in the set of currently generated LSPs from that source in its database (i.e. it is an LSP generated by this Intermediate system), perform the actions described in 7.3.16.1.
 - e) Otherwise, (the source *S* is some other system),
 - 1) If the LSP is newer than the one in the database, or if an LSP from that source does not yet exist in the database:

- i. Store the new LSP in the database, overwriting the existing database LSP for that source (if any) with the received LSP.
 - ii. Set **SRMflag** for that LSP for all circuits other than C.
 - iii. Clear **SRMflag** for C.
 - iv. If C is a non-broadcast circuit, set **SSNflag** for that LSP for C.
 - v. Clear **SSNflag** for that LSP for the circuits associated with a linkage other than C.
- 2) If the LSP is equal to the one in the database (same Sequence Number, Remaining Lifetimes both zero or both non-zero, same checksums):
- i. Clear **SRMflag** for C.
 - ii. If C is a non-broadcast circuit, set **SSNflag** for that LSP for C.
- 3) If the LSP is older than the one in the database:
- i. Set **SRMflag** for C.
 - ii. Clear **SSNflag** for C.

When storing a new LSP, the Intermediate system shall first ensure that it has sufficient memory resources to both store the LSP and generate whatever internal data structures will be required to process the LSP by the Update Process. If these resources are not available the LSP shall be ignored. It shall neither be stored nor acknowledged. When an LSP is ignored for this reason the IS shall enter the "Waiting State" (See 7.3.19).

When attempting to store a new version of an existing LSP (with the same LSPID), which has a length less than or equal to that of the existing LSP, the existing LSP shall be removed from the routing information base and the new LSP stored as a single atomic action. This ensures that such an LSP (which may be carrying the LSP Database Overload indication from an overloaded IS) will never be ignored as a result of a lack of memory resources.

7.3.15.2 Action on receipt of a sequence numbers PDU

When a Sequence Numbers PDU (Complete or Partial, see 7.3.17) is received on circuit C the IS shall perform the following functions:

- a) Perform the following PDU acceptance tests:
 - 1) If the SNP was received over a circuit whose **externalDomain** attribute is "True", the IS shall discard the PDU.
 - 2) If a level 1 Intermediate System receives a level 2 SNP, then the IS shall discard the SNP.
 - 3) If a level 2 Intermediate System receives a level 1 SNP on a circuit with **manualL2OnlyMode** set to "True", then the IS shall discard the SNP.

If circuit C is a broadcast circuit and either

- i. this is a level 1 PSNP and this IS is not the level 1 designated IS for the circuit C, or
 - ii. this is a level 2 PSNP and this IS is not the level 2 designated IS for the circuit C, then
- the IS shall discard the PDU.

- 4) If the **ID Length** field of the PDU is not equal to the value of the IS's **routingDomainIDLength**, the PDU shall be discarded and an **iDFieldLengthMismatch** event generated.
- 5) If this is a level 1 SNP, and the **Maximum Area Addresses** field of the PDU is not equal to the value of the IS's **maximumAreaAddresses** then the PDU shall be discarded and a **maximumAreaAddresses-Mismatch** event generated, unless the IS only implements a value of three for **maximumArea-Addresses**, in which case this check may be omitted.
- 6) If the SNP is received on a broadcast circuit and the source subnetwork address of the SNP does not match the **neighborSNPAddress** of an adjacency of the same level (e.g. a level 1 SNP with a level 1 adjacency) on the circuit over which the SNP was received, then the IS shall discard the SNP without generating an event.

If the SNP is received on a non-broadcast circuit and there is no adjacency of the same level (e.g. a level 1 SNP with a level 1 or level 1 & 2 adjacency) on the circuit over which the SNP was received, then the IS shall discard the SNP without generating an event.

- 7) If this is a level 1 SNP and the value of **area-TransmitPassword** or the set of **areaReceive-Passwords** is non-null, then perform the following tests:
 - i. If the PDU does not contain the **Authentication Information** field then the PDU shall be discarded and an **authenticationFailure** event generated.
 - ii. If the PDU contains the **Authentication Information** field, but the **Authentication Type** is not equal to "Password", then:
 - a) If the IS implements the authentication procedure indicated by the **Authentication Type** whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - b) If the IS does not implement the authentication procedure indicated by the **Authentication Type** then the IS shall ignore the PDU and generate an **authenticationFailure** event.
 - iii. Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of **areaReceivePasswords**, augmented by the value of the **areaTransmitPassword**. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the above values, then the IS shall ignore the PDU and generate an **authentication-Failure** event.
- 8) If this is a level 2 SNP, and the value of **domain-TransmitPassword** or the set of **domainReceive-Passwords** is non-null, then perform the following tests:
 - i. If the PDU does not contain the **Authentication Information** field then the PDU shall be discarded and an **authenticationFailure** event generated.
 - ii. If the PDU contains the **Authentication Information** field, but the **Authentication Type** is not equal to "Password", then:
 - a) If the IS implements the authentication procedure indicated by the **Authentication Type** whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - b) If the IS does not implement the authentication procedure indicated by the **Authentication Type** then the IS shall ignore the PDU and generate an **authenticationFailure** event.
 - iii. Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of **domainReceivePasswords**, augmented by the value of the **domainTransmitPassword**. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the above values, then the IS shall ignore the PDU and generate an **authentication-Failure** event.
- b) For each LSP reported in the Sequence Numbers PDU:
 - 1) Determine if the LSP entry in the received SNP describes an LSP which is newer than a corresponding stored LSP by performing the operations specified in 7.3.16.2 and 7.3.16.3.
 - 2) If the reported value equals the database value and *C* is a non-broadcast circuit, Clear **SRMflag** for *C* for that LSP
 - 3) If the reported value is older than the database value, Clear **SSNflag**, and Set **SRMflag**.
 - 4) If the reported value is newer than the database value, Set **SSNflag**, and if *C* is a non-broadcast circuit Clear **SRMflag**.
 - 5) If no database entry exists for the LSP, and the reported **Remaining Lifetime**, **Checksum** and **Sequence Number** fields of the LSP are all non-zero, create an entry with sequence number 0 (see 7.3.16.1), and set **SSNflag** for that entry and circuit *C*. Under no circumstances shall **SRMflag** be set for such an LSP with zero sequence number.

NOTE 32 This is because possessing a zero sequence number LSP is semantically equivalent to having no information about that LSP. If such LSPs were propagated by setting **SRMflag** it would result in an unnecessary consumption of both bandwidth and memory resources.

- c) If the Sequence Numbers PDU is a Complete Sequence Numbers PDU, Set **SRMflags** for *C* for all LSPs in the database (except those with zero sequence number or zero **Remaining Lifetime**) with LSPIDs within the range specified for the CSNP by the **Start LSPID** and **End LSPID** fields, which were not mentioned in the Complete Sequence Numbers PDU (i.e. LSPs this system has, which the neighbour does not claim to have).

7.3.15.3 Action on expiration of complete SNP interval

The IS shall perform the following actions every **CompleteSNPInterval** for circuit *C*:

- a) If *C* is a broadcast circuit, then

- 1) If this Intermediate system is a Level 1 Designated Intermediate System on circuit C, transmit a complete set of Level 1 Complete Sequence Numbers PDUs on circuit C. Ignore the setting of **SSNflag** on Level 1 Link State PDUs.

If the value of the IS's **areaTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted CSNP, indicating an **Authentication Type** of "Password" and containing the **areaTransmitPassword** as the authentication value.

- 2) If this Intermediate system is a Level 2 Designated Intermediate System on circuit C, transmit a complete set of Level 2 Complete Sequence Numbers PDUs on circuit C. Ignore the setting of **SSNflag** on Level 2 Link State PDUs.

If the value of the IS's **domainTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted CSNP, indicating an **Authentication Type** of "Password" and containing the **domainTransmitPassword** as the authentication value.

A complete set of CSNPs is a set whose **Start LSPID** and **End LSPID** ranges cover the complete possible range of LSPIDs. (i.e. there is no possible LSPID value which does not appear within the range of one of the CSNPs in the set). Where more than one CSNP is transmitted on a broadcast circuit, they shall be separated by an interval of at least **minimumBroadcast-LSPTransmissionInterval**.

NOTE 33 An IS is permitted to transmit a small number of CSNPs (no more than 10) with a shorter separation interval, (or even "back to back"), provided that no more than $1000 / \text{minimumBroadcast-LSPTransmissionInterval}$ CSNPs are transmitted in any one second period.

- b) Otherwise (C is a point-to-point circuit, including non-DA DED circuits and virtual links), do nothing. CSNPs are only transmitted on point-to-point circuits at initialisation.

7.3.15.4 Action on expiration of partial SNP interval

The maximum sized Level 1 or Level 2 PSNP which may be generated by a system is controlled by the values of **originatingL1LSPBufferSize** or **originatingL2LSPBufferSize** respectively. An Intermediate system shall perform the following actions every **partialSNPInterval** for circuit C with jitter applied as described in 10.1:

- a) If C is a broadcast circuit, then

- 1) If this Intermediate system is a Level 1 Intermediate System or a Level 2 Intermediate System with **manualL2OnlyMode** "False", but is **not** a Level 1 Designated Intermediate System on circuit C, transmit a Level 1 Partial Sequence Numbers PDU on circuit C, containing entries for as many Level 1 Link State PDUs with **SSNflag** set as will fit in the PDU, and then clear **SSNflag** for these entries. To avoid the possibility of starvation, the scan of the LSP database for those with **SSNflag** set shall commence with the next LSP which was not included in the previous scan. If there were no Level 1 Link State PDUs with **SSNflag** set, do not transmit a Level 1 Partial Sequence Numbers PDU.

If the value of the IS's **areaTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted PSNP, indicating an **Authentication Type** of "Password" and containing the **areaTransmitPassword** as the authentication value.

- 2) If this Intermediate system is a Level 2 Intermediate System, but is **not** a Level 2 Designated Intermediate System on circuit C, transmit a Level 2 Partial Sequence Numbers PDU on circuit C, containing entries for as many Level 2 Link State PDUs with **SSNflag** set as will fit in the PDU, and then clear **SSNflag** for these entries. To avoid the possibility of starvation, the scan of the LSP database for those with **SSNflag** set shall commence with the next LSP which was not included in the previous scan. If there were no Level 2 Link State PDUs with **SSNflag** set, do not transmit a Level 2 Partial Sequence Numbers PDU.

If the value of the IS's **domainTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted PSNP, indicating an **Authentication Type** of "Password" and containing the **domainTransmitPassword** as the authentication value.

- b) Otherwise (C is a point-to-point circuit, including non-DA DED circuits and virtual links)

- 1) If this system is a Level 1 Intermediate system, transmit a Level 1 Partial Sequence Numbers PDU on circuit C, containing entries for as many Level 1 Link State PDUs with **SSNflag** set as will fit in the PDU, and then clear **SSNflag** for these entries. To avoid the possibility of starvation, the scan of the LSP database for those with **SSNflag** set shall commence with the next LSP which was not included in the previous scan. If there were no Level 1 Link State PDUs with **SSNflag** set, do not transmit a Partial Sequence Numbers PDU.

If the value of the IS's **areaTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted PSNP, indicating an **Authentication Type** of "Password" and containing the **areaTransmitPassword** as the authentication value.

- 2) If this system is a Level 2 Intermediate system, transmit a Level 2 Partial Sequence Numbers PDU on circuit C, containing entries for as many Level 2 Link State PDUs with **SSNflag** set as will fit in the PDU, and then clear **SSNflag** for these entries. To avoid the possibility of starvation, the scan of the LSP data base for those with **SSNflag** set shall commence with the next LSP which was not included in the previous scan. If there were no Level 2 Link State PDUs with **SSNflag** set, do not transmit a Partial Sequence Numbers PDU.

If the value of the IS's **domainTransmitPassword** is non-null, then the IS shall include the **Authentication Information** field in the transmitted PSNP, indicating an **Authentication Type** of "Password" and containing the **domainTransmitPassword** as the authentication value.

7.3.15.5 Action on expiration of the minimumLSPTransmissionInterval

An IS shall perform the following action every **minimumLSPTransmissionInterval** with jitter applied as described in 10.1:

- For all point-to-point circuits C (including non-DA DED circuits and virtual links) transmit all LSPs that have **SRMflag** set on circuit C, but **do not** clear the **SRMflag**. The **SRMflag** will subsequently be cleared by receipt of a Complete or Partial Sequence Numbers PDU.

The interval between two consecutive transmissions of the same LSP shall be at least **minimumLSPTransmissionInterval**. Clearly, this can only be achieved precisely by keeping a separate timer for each LSP. This would be an unwarranted overhead. Any technique which ensures the interval will be between **minimumLSPTransmissionInterval** and $2 \times \text{minimumLSPTransmissionInterval}$ is acceptable.

7.3.15.6 Controlling the rate of transmission on broadcast circuits

The attribute **minimumBroadcastLSPTransmissionInterval** indicates the minimum interval between PDU arrivals which can be processed by the slowest Intermediate System on the LAN.

Setting **SRMflags** on an LSP for a broadcast circuit does not cause the LSP to be transmitted immediately. Instead the Intermediate system shall scan the LSP database every **minimumBroadcastLSPTransmissionInterval** (with jitter applied as described in 10.1), and from the set of LSPs which have **SRMflags** set for this circuit, one LSP shall be chosen at random. This LSP shall be multicast on the circuit, and **SRMflags** cleared.

NOTE 34 In practice it would be very inefficient to scan the whole database at this rate, particularly when only a few LSPs had **SRMflags** set. Implementations may require additional data structures in order to reduce this overhead.

NOTE 35 An IS is permitted to transmit a small number of LSPs (no more than 10) with a shorter separation interval, (or even "back to back"), provided that no more than $1000/\text{minimumBroadcastLSPTransmissionInterval}$ LSPs are transmitted in any one second period.

In addition, the presence of any LSPs which have been received on a particular circuit and are queued awaiting processing shall inhibit transmission of LSPs on that circuit. However, LSPs may be transmitted at a minimum rate of one per second even in the presence of such a queue.

7.3.16 Determining the latest information

The Update Process is responsible for determining, given a received link state PDU, whether that received PDU represents new, old, or duplicate information with respect to what is stored in the database.

It is also responsible for generating the information upon which this determination is based, for assigning a sequence number to its own Link State PDUs upon generation, and for correctly adjusting the **Remaining Lifetime** field upon broadcast of a link state PDU generated originally by any system in the domain.

7.3.16.1 Sequence numbers

The sequence number is a 4 octet unsigned value. Sequence numbers shall increase from zero to (**SequenceModulus** - 1). When a system initialises, it shall start with sequence number one for its own Link State PDUs.¹⁾

¹⁾ The IS starts with 1 rather than 0 so that the value 0 can be reserved to be guaranteed to be less than the sequence number of any actually generated Link State PDU. This is a useful property for Sequence Numbers PDUs.

The sequence numbers the Intermediate system generates for its Link State PDUs with different values for LSP number are independent. The algorithm for choosing the numbers is the same, but operationally the numbers will not be synchronised.

If an Intermediate system R somewhere in the domain has information that the current sequence number for source S is greater than that held by S, R will return to S a Link State PDU for S with R's value for the sequence number. When S receives this LSP it shall change its sequence number to be the next number greater than the new one received, and shall generate a link state PDU.

If an Intermediate system needs to increment its sequence number, but the sequence number is already equal to $\text{SequenceModulus} - 1$, the event `attemptToExceedMaximumSequenceNumber` shall be generated and the IS Network entity shall be disabled for a period of at least $\text{MaxAge} + \text{ZeroAgeLifetime}$, in order to be sure that any versions of this LSP with the high sequence number have expired. When it is re-enabled the IS shall start again with sequence number 1.

7.3.16.2 LSP confusion

It is possible for an LSP generated by a system in a previous incarnation to be alive in the domain and have the same sequence number as the current LSP.

To ensure database consistency among the Intermediate Systems, it is essential to distinguish two such PDUs. This is done efficiently by comparing the checksum in a received LSP or an LSP entry in a received SNP with the checksum of the LSP stored in memory.

If the sequence numbers match, but the checksums do not and the LSP is not in the current set of LSPs generated by the local system, then the system that notices the mismatch shall treat the LSP as if its Remaining Lifetime had expired. It shall store one of the copies of the LSP, with zero written as the Remaining Lifetime, and flood the LSP.

If the LSP is in the current set of LSPs generated by the local system then the IS shall change the LSP's sequence number to be the next number greater than that of the received LSP and regenerate the LSP.

7.3.16.3 Remaining lifetime field

When the source generates a link state PDU, it shall set the Remaining Lifetime to **MaxAge**.

When a system holds the information for some time before successfully transmitting it to a neighbour, that system shall decrement the Remaining Lifetime field according to the holding time. Before transmitting a link state PDU to a neighbour, a system shall decrement the Remaining Lifetime in the PDU being transmitted by at least one, or more than one if the transit time to that neighbour is estimated to be greater than one second. When the Remaining Lifetime field reaches zero, the system shall purge that Link State PDU from its database. In order to keep the Intermediate Systems' databases synchronised, the purging of an LSP due to Remaining Lifetime expiration is synchronised by flooding an expired LSP. See 7.3.16.4.

If the Remaining Lifetime in a received LSP or in an LSP entry in an SNP is zero it shall be processed as described in 7.3.16.4. If the Remaining Lifetime in a received LSP or in an LSP entry in a received SNP is non-zero, but there is an LSP in the database with the same sequence number and zero Remaining Lifetime, the LSP in the database shall be considered most recent. Otherwise, the PDU with the larger sequence number shall be considered the most recent.

If the value of Remaining Lifetime is greater than **MaxAge**, the LSP shall be processed as if there were a checksum error.

7.3.16.4 LSP expiration synchronisation

When the Remaining Lifetime on an LSP in memory becomes zero, the IS shall

- a) set all SRMflags for that LSP, and
- b) retain only the LSP header.
- c) record the time at which the Remaining Lifetime for this LSP became zero. When **ZeroAgeLifetime** has elapsed since the LSP Remaining Lifetime became zero, the LSP header shall be purged from the database.

NOTE 36 A check of the checksum of a zero Remaining Lifetime LSP succeeds even though the data portion is not present.

When a purge of an LSP with non-zero Remaining Lifetime is initiated, the header shall be retained for **MaxAge**.

If an LSP from source S with zero Remaining Lifetime is received on circuit C :

- a) If no LSP from *S* is in memory, then the IS shall
 - 1) send an acknowledgement of the LSP on circuit *C*, but
 - 2) shall not retain the LSP after the acknowledgement has been sent.
- b) If an LSP from *S* is in the database, then
 - 1) If the received LSP is newer than the one in the data-base (i.e. received LSP has higher sequence number, or same sequence number and database LSP has non-zero Remaining Lifetime) the IS shall:
 - i. Store the new LSP in the database, overwriting the existing database LSP for that source (if any) with the received LSP.
 - ii. Set **SRMflag** for that LSP for all circuits other than *C*.
 - iii. Clear the **SRMflag** for that LSP for *C*.
 - iv. If *C* is a non-broadcast circuit, set **SSNflag** for that LSP for *C*.
 - v. clear **SSNflag** for that LSP for all circuits associated with that linkage other than *C*.
 - 2) If the received LSP is equal to the one in the database (i.e. same Sequence Number, Remaining Lifetimes both zero) the IS shall
 - i. clear **SRMflag** for *C*, and
 - ii. if *C* is a non-broadcast circuit, set **SSNflag** for that LSP for *C*
 - 3) If the received LSP is older than the one in the data-base (i.e. received LSP has lower sequence number) the IS shall
 - i. set **SRMflag** for *C*, and
 - ii. clear **SSNflag** for *C*.
- c) If this system (or pseudonode) is *S* and there is an un-expired LSP from *S* (i.e. its own LSP) in memory, then the IS
 - 1) shall not overwrite with the received LSP, but
 - 2) shall change the sequence number of the unexpired LSP from *S* as described in 7.3.16.1,
 - 3) generate a new LSP; and
 - 4) set **SRMflag** on all circuits.

7.3.17 Making the update reliable

The update process is responsible for making sure the latest link state PDUs reach every reachable Intermediate System in the domain.

On point-to-point links the Intermediate system shall send an explicit acknowledgement encoded as a Partial Sequence Numbers PDU (PSNP) containing the following information:

- a) source's ID
- b) PDU type (Level 1 or 2)
- c) sequence number
- d) Remaining Lifetime
- e) checksum

This shall be done for all received link state PDUs which are newer than the one in the database, or duplicates of the one in the database. Link state PDUs which are older than that stored in the database are answered instead by a newer link state PDU, as specified in 7.3.14 above.

On broadcast links, instead of explicit acknowledgements for each link state PDU by each Intermediate system, a special PDU known as a Complete Sequence Numbers PDU (CSNP), shall be multicast periodically by the Designated Intermediate System. The PDU shall contain a list of all LSPs in the database, together with enough information so that Intermediate systems receiving the CSNP can compare with their LSP database to determine whether they and the CSNP transmitter have synchronised LSP databases. The maximum sized Level 1 or Level 2 Sequence Numbers PDU which may be generated by a system is controlled by the values of **originatingL1LSPBufferSize** or **originatingL2LSPBufferSize** respectively. In practice, the information required to be transmitted in a single CSNP may be greater than will fit in a single PDU. Therefore

each CSNP carries an inclusive range of LSPIDs to which it refers. The complete set of information shall be conveyed by transmitting a series of individual CSNPs, each referring to a subset of the complete range. The ranges of the complete set of CSNPs shall be contiguous (though not necessarily transmitted in order) and shall cover the entire range of possible LSPIDs.

The LAN Level 1 Designated Intermediate System shall periodically multicast complete sets of Level 1 CSNPs to the multi-destination address AIL1ISS. The LAN Level 2 Designated Intermediate System shall periodically multicast complete sets of Level 2 CSNPs to the multi-destination address AIL2ISS.

Absence of an LSPID from a Complete Sequence Numbers PDU whose range includes that LSPID indicates total lack of information about that LSPID.

If an Intermediate system, upon receipt of a Complete Sequence Numbers PDU, detects that the transmitter was out of date, the receiver shall multicast the missing information.

NOTE 37 Receipt of a link state PDU on a link is the same as successfully transmitting the Link State PDU on that link, so once the first Intermediate system responds, no others will, unless they have already transmitted replies.

If an Intermediate system detects that the transmitter had more up to date information, the receiving Intermediate system shall multicast a Partial Sequence Numbers PDU (PSNP), containing information about LSPs for which it has older information. This serves as an implicit request for the missing information. Although the PSNP is multicast, only the Designated Intermediate System of the appropriate level shall respond to the PSNP.

NOTE 38 This is equivalent to the PSNP being transmitted directly to the Designated Intermediate System, in that it avoids each Intermediate System unnecessarily sending the same LSP(s) in response. However, it has the advantage of preserving the property that all routing messages can be received on the multi-destination addresses, and hence by a LAN adapter dedicated to the multi-destination address.

When a point-to-point circuit (including non-DA DED circuits and virtual links) starts (or restarts), the IS shall

- a) set SRMflag for that circuit on all LSPs, and
- b) send a Complete set of Complete Sequence Numbers PDUs on that circuit.

7.3.18 Validation of databases

An Intermediate System shall not continue to operate for an extended period with corrupted routing information. The IS shall therefore operate in a *fail-stop* manner. If a failure is detected, the Intermediate system Network entity shall be disabled until the failure is corrected. In the absence of an implementation-specific method for ensuring this, the IS shall perform the following at least every maximumLSPGenerationInterval:

- a) On expiration of this timer the IS shall re-check the checksum of every LSP in the LSP database (except those with a Remaining Lifetime of zero) in order to detect corruption of the LSP while in memory. If the checksum of any LSP is incorrect, the event corruptedLSPDetected shall be logged, and as a minimum the entire Link State Database shall be deleted and action taken to cause it to be re-acquired. One way to achieve this is to disable and re-enable the IS Network entity.

NOTE 39 On point-to-point links, this requires at least that a CSNP be transmitted

- b) On completion of these checks the decision process shall be notified of an event (even if any newly generated LSPs have identical contents to the previous ones). This causes the decision process to be run and the forwarding databases re-computed, thus protecting against possible corruption of the forwarding databases in memory, which would not otherwise be detected in a stable topology.
- c) The IS shall reset the timer for a period of maximum-LSPGenerationInterval with jitter applied as described in 10.1.

7.3.19 LSP database overload

As a result of network mis-configuration, or certain transitory conditions, it is possible that there may be insufficient memory resources available to store a received Link State PDU. When this occurs, an IS needs to take certain steps to ensure that if its LSP database becomes inconsistent with the other ISs', that these ISs do not rely on forwarding paths through the over loaded IS.

7.3.19.1 Entering the waiting state

When an LSP cannot be stored, the LSP shall be ignored and Waiting State shall be entered. A timer shall be started for **waitingTime**, and the Intermediate System shall generate and flood its own LSP with zero LSP number with the LSP Database Overload Bit set. This prevents this Intermediate system from being considered as a forwarding path by other Intermediate Systems.

It is possible that although there are sufficient resources to store an LSP and permit the operation of the Update Process on that LSP, the Decision Process may subsequently require further resources in order to complete. If these resources are not available, the Intermediate system shall then (i.e. during the attempt to run the Decision Process) enter Waiting State until such time as they are available and **waitingTime** have elapsed since the last LSP was ignored by the Update Process.

An implementation shall partition the available memory resources between the Level 1 and Level 2 databases. An overload condition can therefore exist independently for Level 1 or Level 2 (or both). The status attributes **I1State** and **I2State** indicate the condition for the Level 1 and Level 2 databases respectively. On entering Level 1 "Waiting State" the IS shall generate the **ISPL1DatabaseOverload** event, and on entering Level 2 "Waiting State" the IS shall generate the **ISPL2DatabaseOverload** event.

7.3.19.2 Actions in level 1 waiting state

While in Level 1 "waiting" state

- a) If a Link State PDU cannot be stored, the IS shall ignore it and restart the timer for **waitingTime**.
- b) The IS shall continue to run the Decision and Forwarding processes as normal.
- c) When the **waitingTime** timer expires, the IS shall:
 - 1) Generate an **ISPL1DatabaseOverload** (recovered) event.
 - 2) Clear the LSP Database Overload bit in its own Level 1 LSP with zero LSP number and re-issue it.
 - 3) Set the **I1State** to "On".
 - 4) Resume normal operation.

7.3.19.3 Actions in level 2 waiting state

While in Level 2 "waiting" state

- a) If a Link State PDU cannot be stored, the IS shall ignore it and restart the timer for **waitingTime** seconds.
- b) The IS shall continue to run the Decision and Forwarding processes as normal.
- c) When the **waitingTime** timer expires, the IS shall:
 - 1) Generate an **ISPL2DatabaseOverload** (recovered) event.
 - 2) Clear the LSP Database Overload bit in its own Level 2 LSP with zero LSP number and re-issue it.
 - 3) Set the **I2State** to "On".
 - 4) Resume normal operation.

7.3.20 Use of the link state database

The only portion of the database relevant to the Decision Process is the data portion of the Link State PDUs. The Update Process additionally uses the fields **Sequence Number**, **Remaining Lifetime**, and variable **SRMflag**. The Remaining Lifetimes in the stored link state PDUs can either be periodically decremented, or converted upon receipt into an internal timestamp, and converted back into a Remaining Lifetime upon transmission.

7.3.20.1 Synchronisation with the decision process

Since the Update process and the Decision process share the link state database, care must be taken that the Update process does not modify the link state database while the Decision process is running.

There are two approaches to this. In one approach, the Decision process signals when it is running. During this time, the Update process queues incoming Link State PDUs, and does not write them into the link state database. If more Link State PDUs arrive than can fit into the queue allotted while the Decision process is running, the Update process drops them and does not acknowledge them.

Another approach is to have two copies of the link state database — one in which the Decision process is computing, and the other in which the Update process initially copies over the first database, and in which all new Link State PDUs are written. Additionally, depending on the hashing scheme, it is likely that a second copy of the address hash table will be required, so that the Update process can do a rehash occasionally for efficiency.

When the Decision process is ready to run again, it locks the new copy of the link state database, leaving the Update process to copy over the information into the first area, and write new updates while the Decision process runs again. The advantage of the first approach is that it takes less memory. The advantage of the second approach is that Link State PDUs will never need to be dropped.

NOTE 40 If the decision process is implemented according to the specification in C.2, a finer level of parallelism is possible, as described below.

Arrival of a Link State PDU for a system before that system has been put into TENT is permitted. The new Link State PDU is used when that system is eventually put into TENT. Similarly, arrival of a new Link State PDU for a system after that system has been put into PATHS is permitted. That system has already been completely processed. The arrival of the new Link State PDU is noted and the decision process re-executed when the current execution has completed. An in-progress execution of the decision process shall not be abandoned, since this could prevent the decision process from ever completing.

Arrival of a Link State PDU for a system between that system being put on TENT and being transferred to PATHS can be treated as equivalent to one of the previous two cases (for example, by buffering, or taking some corrective action).

7.3.20.2 Use of buffers and link bandwidth

Implementations shall have a buffer management strategy that does not prevent other clients of the buffering service from acquiring buffers due to excessive use by the Update Process. They shall also ensure that the Update Process does not consume all the available bandwidth of links. In particular no type of traffic should experience starvation for longer than its acceptable latency. Acceptable latencies are approximately as follows:

- Hello traffic – Hello timer $\times 0,5$
- Data Traffic – 10 s.

NOTE 41 The first of these requirements can be met by restricting the Update process to the use of a single buffer on each circuit for transmission. This may also cause the second requirement to be met, depending on the processor speed.

7.3.21 Parameters

MaxAge – This is the amount of time that may elapse since the estimated origination of the stored Link State PDU by the source before the LSP is considered expired. The expired LSP can be deleted from the database after a further **ZeroAgeLifetime** has expired. **MaxAge** shall be larger than **maximumLSPGenerationInterval**, so that a system is not purged merely because of lack of events for reporting Link State PDUs.

MaxAge is an architectural constant equal to 20 min.

ZeroAgeLifetime — This is the minimum amount of time for which the header of an expired LSP shall be retained after it has been flooded with zero Remaining Lifetime. A very safe value for this would be $2 \times \text{MaxAge}$. However all that is required is that the header be retained until the zero Remaining Lifetime LSP has been safely propagated to all the neighbours.

ZeroAgeLifetime is an architectural constant with a value of 1 min.

maximumLSPGenerationInterval – This is the maximum amount of time allowed to elapse between generation of Link State PDUs by a source. It shall be less than **MaxAge**.

Setting this parameter too fast adds overhead to the algorithms (a lot of Link State PDUs). Setting this parameter too slow (and not violating constraints) causes the algorithm to wait a long time to recover in the unlikely event that incorrect Link State information exists somewhere in the domain about the system.

A reasonable setting is 15 min.

minimumLSPGenerationInterval – This is the minimum time interval between generation of Link State PDUs. A source Intermediate system shall wait at least this long before re-generating one of its own Link State PDUs.

Setting the interval too slow causes a delay in reporting new information. Setting the interval too fast allows too much overhead.

A reasonable setting is 30 s.

minimumLSPTransmissionInterval – This is the amount of time an Intermediate system shall wait before further propagating another Link State PDU from the same source system.

Setting the interval too slow causes a delay in propagation of routeing information and stabilisation of the routeing algorithm. Setting the interval too fast allows the possibility that the routeing algorithm, under low probability circumstances, will use too many resources (CPU and bandwidth).

Setting **minimumLSPTransmissionInterval** greater than **minimumLSPGenerationInterval** makes no sense, because the source would be allowed to generate LSPs more quickly than they would be allowed to be broadcast. Setting **minimumLSPTransmissionInterval** smaller than **minimumLSPGenerationInterval** is desirable to recover from lost LSPs.

A reasonable value is 5 s.

CompleteSNPIInterval – This is the amount of time between periodic transmissions of a complete set of Sequence Number PDUs by the Designated Intermediate system on a broadcast link. Setting the interval too slow delays convergence of the routeing algorithm when Link State PDUs are lost due to the datagram environment of the Data Link layer on the broadcast link.

Setting the interval too fast results in extra control traffic overhead.

A reasonable value is 10 s.

partialSNPIInterval – This is the amount of time between periodic action for transmission of Partial Sequence Number PDUs. It shall be less than **minimumLSPTransmissionInterval**.

Setting the interval too slow delays convergence of the routeing algorithm when Link State PDUs are lost due to the environment of the Data Link layer on the corresponding link. Setting the interval too fast results in extra control traffic overhead.

A reasonable value is 2 s.

7.4 Forwarding process

The Forwarding process is responsible both for transmitting NPDU's originated by this system, and for forwarding NPDU's originated by other systems

7.4.1 Input and output

INPUT

- NPDU's from the ISO 8473 protocol machine
- PDU's from Update Process
- PDU's from Receive Process
- Forwarding Databases (Level 1 and 2) — one for each routeing metric

OUTPUT

- PDU's to Data Link Layer

7.4.2 Routeing metric selection

The Forwarding process selects a forwarding database for each NPDU to be relayed based on:

- the level at which the forwarding is to occur: level 1 or level 2; and
- a mapping of the ISO 8473 QoS Maintenance field onto one of the Intermediate system's supported routeing metrics.

The former selection is made by examining the Destination Address field of the NPDU. The latter selection is made as follows:

- a) If the QoS Maintenance field is not present in the NPDU, then the IS shall select the forwarding database calculated for the *default metric*.
- b) If the QoS Maintenance field is present, the IS shall examine bits 8 and 7 of the parameter value octet (QoS format code). If the QoS format code represents any other than “00” (meaning globally unique with strong forwarding) and “11” (meaning globally unique with weak forwarding), then the IS shall select the forwarding database calculated for the default metric, otherwise
- c) The IS shall select a forwarding database by mapping the values of bit 3, 2 and 1 of the parameter value shown in Table 1 and shall proceed as follows;
 - 1) If the QoS format code represents “00”, the IS shall forward based on the selected optional metric.
 - 2) If the QoS format code represents “11”, the IS shall proceed as follows;
 - i. If the IS does not support the selected routing metric, the IS shall forward based on the default metric;
 - ii. If the forwarding database for one of the optional routing metrics is selected and the database either does not contain an entity for the Destination Address in the NPDU being relayed, or contains an entry indicating that the destination is unreachable using that metric, then the IS shall attempt to forward based upon the default metric;
 - iii. Otherwise, forward based on the selected optional metric.

Table 1 – QoS Maintenance bits to routing metric mappings

bit 1	bit 2	bit 3	Selected Routing Metric
0	0	0	expense metric
1	0	0	default metric
0	1	1	default metric
0	1	0	expense metric
0	0	1	delay metric
1	1	0	error metric
1	0	1	delay metric
1	1	1	error metric

7.4.3 Forwarding decision

7.4.3.1 Basic operation

Let DEST = the Network Layer destination address of the PDU to be forwarded, or the next entry in the source routing field, if present. It consists of sub-fields Area Address, ID, and SEL.

NOTE 42 The SEL field in the destination address is not examined by Intermediate Systems. It is used by End Systems to select the proper Transport entity to which to deliver NSDUs.

This system's (the one examining this PDU for proper forwarding decision) address consists of sub-fields area address and ID.

- a) If the local system type is a level 1 Intermediate system, **or** the local system type is a level 2 Intermediate system and $\text{AttachedFlag}_k = \text{False}$, then:
 - 1) If the Area Address in the PDU to be forwarded matches any one of the area addresses of this IS, then consult the level 1 forwarding database to determine the adjacency which is the next hop on the path to the NPDU's destination. Forward the NPDU on this adjacency.
 - 2) Otherwise, consult the level 1 forwarding database to determine the adjacency which is the next hop on the path to the nearest level 2 IS in the area, and forward the NPDU on this adjacency.
- b) If the local system type is of iSType “L2 Intermediate System” and AttachedFlag_k is “True” (where k means the metric selected based on the above procedure) then:
 - 1) If the Area Address of the DEST matches any one of the area addresses of this area,
 - i. If this PDU is a Data PDU of which both SEL fields contain the IntradomainRoutingSelector value and k means default metric (i.e., this PDU was encapsulated by a Designated partition level 2 IS), consult the level 2 Forwarding database to determine the next hop level 2 IS on the path to the partition Designated level 2 IS. Forward the NPDU on this adjacency

NOTE: A level 2 route shall be preferred over a level 1 route that traverses a virtual link.

- ii. Otherwise, consult the level 1 forwarding database to determine the adjacency which is the next hop on the path to the NPDU's destination. Forward the NPDU on this adjacency.
- 2) If the Area Address of the DEST does not match any area address of this area, consult the level 2 forwarding database to determine the adjacency which is the next hop on the path to the destination area, and forward the NPDU on this adjacency. In case of multiple matching address prefixes, the path with the longest prefix shall be chosen.

7.4.3.2 Encapsulation for partition repair

If this Intermediate system is the Partition Designated Level 2 IS for this partition, and the PDU is being forwarded onto the special adjacency to a Partition Designated Level 2 Intermediate system in a different partition of this area, encapsulate the complete PDU as the data field of a data NPDU (i.e., with an additional layer of header), making this system the Source address and the other Partition Designated Level 2 Intermediate system (obtained from the identifier attribute of the Virtual Adjacency managed object) the Destination Address field in the outer PDU header. Set the header of the outer PDU to indicate forwarding via the default routeing metric [see 7.2.10.4, item d)]. Then forward the encapsulated PDU as described in 7.4.3.3 below.

7.4.3.3 The procedure forward

Choose either the level 1 or level 2 forwarding database, depending on the destination Network Address in the NPDU. From that database, select the adjacency for the next hop to that destination. If forwarding at level 1 for a destination which is not in the area, choose the adjacency for the nearest Level 2 IS computed as described in 7.2.9.1.

If there are multiple possible adjacencies, as a result of multiple minimum cost paths, then one of those adjacencies shall be chosen. An implementation may choose the adjacency at random, or may use the possible adjacencies in "round robin" fashion.

If there is no entry in the selected forwarding database for the destination Network Address, and the NPDU originated from the local Transport entity and the system has one or more Intermediate System adjacencies, then one of those is chosen at random (or in "round robin" fashion). Otherwise the procedure returns the value "False".¹⁾

NOTE 43 Since the local adjacency database is pre-loaded into the decision process, there will always be an entry in the forwarding database for destinations to which an adjacency exists.

NOTE 44 The PDU to be forwarded may require fragmentation, depending on which circuit it is to be forwarded over.

7.4.3.4 Generating redirect PDUs

In addition to forwarding an NPDU, the IS shall inform the local ISO 9542 protocol machine to generate a *Redirect PDU* if the PDU is being forwarded onto the same circuit from which it came, and if the source SNPA address of the NPDU indicates that the NPDU was received from an End System.

7.4.4 Receive process

The Receive process is passed information from any of the following sources.

- received PDUs with the NLPID of IntraDomain-RouteingPD,
- configuration information from the ISO 9542 protocol machine,
- ISO 8473 data PDUs handed to the routeing function by the ISO 8473 protocol machine.

When an area is partitioned, a level 2 path is used as a level 1 link to repair the partitioned area. When this occurs, all PDUs (between the neighbours which must utilise a multi-hop path for communication) shall be encapsulated in a data NPDU, addressed to the IntraDomainRouteingSelector. Control traffic (LSPs, Sequence Numbers PDUs) shall also be encapsulated, as well as data NPDUs that are to be passed between the "neighbours".

¹⁾ This is done so that a system in the overloaded state will still be able to originate or forward NPDUs. If a system with a partial routeing information base were prohibited from attempting to forward to an unknown destination, system management would be unable to either communicate with this system, or route through it, for the purpose of diagnosing and/or correcting the underlying fault.

NOTE 45 It is not necessary to transmit encapsulated IIH PDUs over a virtual link, since virtual adjacencies are established and monitored by the operation of the Decision Process and not the Subnetwork Dependent functions.

7.4.4.1 Basic operation

The Receive Process shall perform the following functions on each received PDU:

- If it is a Link State PDU, pass it to the Update Process
- If it is a Sequence Numbers PDU, pass it to the Update Process
- If it is an IIH PDU, pass it to the appropriate Subnetwork Dependent Function
- If it is a data NPDU or Error Report for another destination, pass it to the Forwarding Process
- Otherwise, ignore the PDU

7.4.4.2 Decapsulation

If an ISO 8473 data NPDU, addressed to *this system*, is received, and the SEL field of the address is equal to *IntraDomainRouteingSelector*, then the IS shall

- decapsulate the NPDU (remove the outer NPDU header).
- If the decapsulated PDU is an ISO 8473 NPDU, move the “congestion” indications to the decapsulated NPDU, and pass it to the ISO 8473 protocol machine.
- Otherwise, if the decapsulated PDU is not an ISO 8473 PDU, perform the following steps on the decapsulated PDU:
 - If it is a link state PDU or Sequence Numbers PDU, pass it to the Update process;
 - Otherwise, ignore the PDU.

If an ISO 8473-1 Error NPDU, addressed to this system, is received, and the SEL field of the address is equal to *IntraDomainRouteingSelector*, then the IS shall notify this event to System Management and discard this PDU.

7.5 Routeing constants and parameters

The architectural constants are described in table 2.

The routeing parameters settable by System Management are listed for each managed object in clause 11.

Table 2 – Routeing architectural constants

Name	Value	Description
MaxLinkMetric	63	Maximum value of a routeing metric assignable to a circuit
MaxPathMetric	1023	Maximum total metric value for a complete path
ISO-SAP	FE	The SAP for ISO Network Layer on ISO 8802-2 LANs. This SAP value is used for transmission and reception of all PDUs of this protocol.
IntradomainRouteingPD	10000011	The Network Layer Protocol Identifier assigned to this protocol, as recorded in ISO/TR 9577
IntradomainRouteingSelector	0	The NSAP selector for the Intermediate System Network entity
Sequence Modulus	2^{32}	Size of the sequence number space used by the Update Process
RecieveLSPBufferSize	1492	The size of LSP which all Intermediate systems must be capable of recieving
MaxAge	1200	Number of seconds before LSP considered expired
ZeroAgeLifetime	60	Number of seconds that an LSP with zero Remaining Lifetime shall be retained after propagating a purge.
ISISHoldingMultiplier	10	The number by which to multiply iSISHelloTimer to obtain Holding Timer for Level 1 and Level 2 IIH PDUs.
Jitter	25%	The percentage of jitter which is applied to the generation of periodic PDUs. See 10.1 for further information on generating jitter on timers.

8 Subnetwork dependent functions

The *Subnetwork Dependent Functions* mask the characteristics of the different kinds of Subnetworks from the *Subnetwork Independent Routing Functions*. The only two types of circuits the Subnetwork Independent Functions recognise are *broadcast* and *general topology*.

The Subnetwork Dependent Functions include:

- The use of the ISO 8473 *Subnetwork Dependent Convergence Functions (SND CF)* so that this protocol may transmit and receive PDUs over the same subnetwork types, using the same techniques, as does ISO 8473.
- Co-ordination with the operation of the ES-IS protocol (ISO 9542) in order to determine the Network layer addresses (and on Broadcast subnetworks, the subnetwork point of attachment address) and identities (End System or Intermediate System) of all adjacent neighbours. This information is held in the *Adjacency* database. It is used to construct Link State PDUs.
- The exchange of IIH PDUs. While it is possible for an Intermediate System to identify that it has an Intermediate System neighbour by the receipt of an ISO 9542 ISH PDU, there is no provision within ISO 9542 to indicate whether the neighbour is a Level 1 or a Level 2 Intermediate System. Specific PDUs (*LAN Level 1*, *LAN Level 2* and *Point-to-point IIH PDUs*) are defined to convey this information.

8.1 Multi-destination circuits on ISs at a domain boundary

Routeing information (e.g. Link State PDUs) is not exchanged across a routeing domain boundary. All routeing information relating to a circuit connected to another routeing domain is therefore entered via the Reachable Address managed objects. This information is disseminated to the rest of the routeing domain via Link State PDUs as described in 7.3.3.2. This has the effect of causing NPDUs destined for NSAPs which are included in the **addressPrefix** of the Reachable Addresses to be relayed to that Intermediate System at the domain boundary. On receipt of such an NPDUs the Intermediate system shall forward it onto the appropriate circuit, based on its own Link State information. However in the case of multi-destination subnetworks (such as an ISO 8208 subnetwork using Dynamic Assignment, a broadcast subnetwork, or a connectionless subnetwork) it is necessary to ascertain additional subnetwork dependent addressing information in order to forward the NPDUs to a suitable SNPA. (This may be the target End system or an Intermediate system within the other domain.)

In general the SNPA address to which an NPDUs is to be forwarded can be derived from the destination NSAP of the NPDUs. It may be possible to perform some algorithmic manipulation of the NSAP address in order to derive the SNPA address. However there may be some NSAPs where this is not possible. In these cases it is necessary to have pre-configured information relating an address prefix to a particular SNPA address.

This is achieved by additional information contained in the reachable address managed object. The **mappingType** attribute specifies the means by which next hop subnetwork addressing information can be derived for NPDUs forwarded based upon a given address prefix. The **mappingType** attribute may be specified as:

explicit — The SNPA address or set of SNPA addresses is manually pre-configured as an attribute of the reachable address managed object.

extractIDI — The SNPA is embedded in the IDI of the destination NSAP address according to the format and encoding rules of ISO 8348. This SNPA extraction algorithm can be used in conjunction with destination addresses from the X.121, F.69, E.163, and E.164 addressing subdomains.

extractDSP — All or a suffix of the SNPA is embedded in the DSP of the destination address. This SNPA extraction algorithm requires manual pre-configuration of **sNPAMask** and **sNPAPrefix** attributes of the reachable address managed object. The **sNPAMask** attribute is a bit mask with 1s indicating the location of the SNPA (suffix) within the destination NSAP DSP. The part of the SNPA extracted from the NSAP is appended to the **sNPAPrefix** to form the next hop subnetwork addressing information.

An example of a set of Reachable Addresses is shown in table 3.

The table is interpreted as follows:

- a) For the ISO DCC prefix 39 123, use the SNPA address X.

- b) For the X.121 IDI address prefix 37 aaaaa, do not use aaaaa, but use *B* instead.
- c) For all IDPs based on SNPAs with DNIC *D* (i.e. with address prefix 37 *D*), use the address *Y* (which would probably be a gateway to a subnetwork with DNIC *D*).
- d) For any other X.121 IDI (i.e. address prefix 37) – use the SNPA whose address is used as the IDI.
- e) For the ISO ICD prefix 47 0005 C0 use the SNPA address formed by concatenating *Z* with next 6 octets of the DSP following the 47 0005 C0 prefix.
- f) Anything else (“*” in table 3) – use one of the SNPA addresses *R*, *S* or *T*. These would typically be the SNPA addresses of Level 2 Intermediate Systems through which any other addresses could potentially be reached.

8.2 Point-to-point subnetworks

This clause describes the identification of neighbours on both point-to-point links and Static circuits.

The IS shall operate the ISO 9542 protocol, shall be able to receive ISO 9542 ISH PDUs from other ISs, and shall store the information so obtained in the adjacency database.

8.2.1 Receipt of ESH PDUs — database of end systems

An IS shall enter an End system into the adjacency database when an ESH PDU is received on a circuit. If an ESH PDU is received on the same circuit, but with a different NSAP address, the new address shall be added to the adjacency, with a separate timer. A single ESH PDU may contain more than one NSAP address. When a new data link address or NSAP address is added to the adjacency database, the IS shall generate an **adjacencyStateChange** (Up) event on that adjacency.

The IS shall set a timer for the value of **Holding Time** in the received ESH PDU. If another ESH PDU is not received from the ES before that timer expires, the ES shall be purged from the database, provided that the Subnetwork Independent Functions associated with initialising the adjacency have been completed. Otherwise the IS shall clear the adjacency as soon as those functions are completed.

Table 3 – Example of reachable address information

Address Prefix	Mapping Type	SNPA Address
39 123	Explicit	X
37 aaaa	Explicit	B
37 D	Explicit	Y
37	extractIDI	Extract X.121 SNPA addrss from NSAP IDI
47 00055 C0	extractDSP	sNPAPrefix=Z sNPAMask=00000000FFFFFFFFFFFF
*	Explicit	R,S,T

When the adjacency is cleared, the Subnetwork Independent Functions shall be informed of an **adjacencyStateChange** (Down) event, and the adjacency can be reused after the Sub network Independent Functions associated with bringing down the adjacency have been completed.

8.2.2 Receiving ISH PDUs by an intermediate system

On receipt of an ISH PDU by an Intermediate System, the IS shall create an adjacency (with **adjacencyState** “Initialising” and **neighbourSystemType** “Unknown”), if one does not already exist, and then perform the following actions:

- a) If the **adjacencyState** is “Up” and the ID portion of the NET field in the ISH PDU does not match the **neigh-bourID** of the adjacency then the IS shall
 - 1) generate an **adjacencyStateChange** (Down) event;
 - 2) delete the adjacency; and
 - 3) create a new adjacency with:
 - i. **adjacencyState** set to “Initialising”, and
 - ii. **neighbourSystemType** set to “Unknown”.
 - 4) perform the following actions.

- b) If the **adjacencyState** is “Initialising”, and the **neighbourSystemType** status is “Intermediate System”, the ISH PDU shall be ignored.
- c) If the **adjacencyState** is “Initialising” and the **neighbourSystemType** status is not “Intermediate System”, a point-to-point IIH PDU shall be transmitted as described in 8.2.4.
- d) The **neighbourSystemType** shall be set to “Intermediate System” indicating that the neighbour is an Intermediate system, but the type (L1 or L2) is, as yet, unknown.
- e) The **circuitType** field shall be set according to Table 4.

8.2.3 Sending ISH PDUs by an Intermediate System

An Intermediate System shall cause ISO 9542 to send ISH PDUs whenever a point-to-point circuit is either first enabled or the adjacency on the circuit transitions to the Down state. ISH PDUs shall continue to be sent periodically at an interval specified by the ISO 9542 Intermediate System Configuration Timer until an adjacency on the circuit transitions to the Up state.

ISH PDUs are not required to be sent when an adjacency on the circuit is in the Up state and it is suggested that an implementation not send them in this state. ISH PDUs received on a point-to-point circuit when an adjacency is in the Up state should be ignored.

8.2.4 Sending point-to-point IIH PDUs

An IS shall send Point-to-Point IIH PDUs on those Point-to-Point circuits whose **externalDomain** attribute is set “False”. The IIH PDU shall be sent when:

- a) the IS receives an ISH PDU
- b) whenever **iSISHelloTimer** expires

The **iSISHelloTimer** shall be (re)started upon transmission of the IIH PDU.

The IIH shall be constructed and transmitted as follows:

- a) The **Circuit Type** field shall be set according to table 4.
- b) The **Local Circuit ID** field shall be set to a value assigned by this Intermediate system when the circuit is created. This value shall be unique among all the circuits of this Intermediate system.
- c) The first Point-to-point IIH PDU (i.e. that transmitted as a result of receiving an ISH PDU, rather than as a result of timer expiration) shall be padded (with trailing **PAD** option fields containing arbitrary valued octets) so that the SNSDU containing the IIH PDU has a length of at least *maxsize* - 1 octets¹⁾ where *maxsize* is the maximum of
 - 1) **dataLinkBlockSize**
 - 2) **originatingL1LSPBufferSize**
 - 3) **originatingL2LSPBufferSize**

This is done to ensure that an adjacency will only be formed between systems which are capable of exchanging PDUs of length up to *maxsize* octets. In the absence of this check, it would be possible for an adjacency to exist with a lower maximum block size, with the result that some LSPs and SNPs (i.e. those longer than this maximum, but less than *maxsize*) would not be exchanged.

NOTE 46 It is necessary for the manager to ensure that the value of **dataLinkBlockSize** on a circuit which will be used to form an Intermediate system to Intermediate system adjacency is set to a value greater than or equal to the maximum of the **LSPBufferSize** characteristics listed above. If this is not done, the adjacency will fail to initialise. It is not possible to enforce this requirement, since it is not known until initialisation time whether or not the neighbour on the circuit will be an End system or an Intermediate system. An End system adjacency may operate with a lower value for **dataLinkBlockSize**.

- d) If the value of the **circuitTransmitPassword** for the circuit is non-null, then the IS shall include the **Authentication Information** field in the transmitted IIH PDU, indicating an **Authentication Type** of “Password” and containing the **circuitTransmitPassword** as the authentication value.
 - 1) The minimum length of **PAD** which may be added is 2 octets, since that is the size of the option header. Where possible the PDU should be padded to *maxsize*, but if the PDU length is *maxsize* - 1 octets no padding is possible (or required).

Table 4 – Setting the value of the circuit type field

iSType	Circuit manualL2OnlyMode	Circuit Type Field
Level 1	--	Level 1 only (1)
Level 2	"True"	Level 2 only (2)
Level 2	"False"	Level 1 and 2 (3)

8.2.5 Receiving point-to-point IIH PDUs

8.2.5.1 PDU acceptance tests

On receipt of a Point-to-Point IIH PDU, perform the following PDU acceptance tests:

- If the IIH PDU was received over a circuit whose **externalDomain** attribute is set "True", the IS shall discard the PDU.
- If the **ID Length** field of the PDU is not equal to the value of the IS's **routingDomainIDLength**, the PDU shall be discarded and an **IDFieldLengthMismatch** event generated.
- If the value of **circuitTransmitPassword** or the set of **circuitReceivePasswords** for this circuit is non-null, then perform the following tests:
 - If the PDU does not contain the **Authentication Information** field then the PDU shall be discarded and an **authenticationFailure** event generated.
 - "If the PDU contains the **Authentication Information** field, but the **Authentication Type** is not equal to "Password", then:
 - If the IS implements the authentication procedure indicated by the **Authentication Type** whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - If the IS does not implement the authentication procedure indicated by the **Authentication Type** then the IS shall ignore the PDU and generate an **authenticationFailure** event."
 - Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of **circuitReceivePasswords** for the circuit on which the PDU was received. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the **circuit-ReceivePasswords**, then the IS shall ignore the PDU and generate an **authenticationFailure** event.

8.2.5.2 IIH PDU Processing

When a Point-to-point IIH PDU is received by an Intermediate system, the area addresses of the two Intermediate Systems shall be compared to ascertain the validity of the adjacency. If the two Intermediate systems have an area address in common and matching values for **maximumAreaAddresses**, the adjacency is valid for all combinations of Intermediate system types (except where a Level 1 Intermediate system is connected to a Level 2 Intermediate system with **manualL2OnlyMode** set "True"). However, if they have no area address in common, the adjacency is only valid if both Intermediate systems are Level 2, and the IS shall mark the adjacency as Level 2 Only. This is described in more detail below.

On receipt of a Point-to-point IIH PDU, each of the **Area Addresses** from the PDU shall be compared with the set of area addresses in the **manualAreaAddresses** attribute.

- If a match is detected between any pair the following actions are taken.
 - If the **Maximum Area Addresses** field of the PDU is not equal to the value of the IS's **maximumAreaAddresses** then the PDU shall be discarded and a **maximumAreaAddresses-Mismatch** event generated, unless the IS only implements a value of three for **maximumArea-Addresses**, in which case this check may be omitted.
 - If the local system is of iSType "L1IntermediateSystem" the IS shall perform the action indicated by table 5.
 - If the local system is of iSType "L2Intermediate-System" and the Circuit **manualL2OnlyMode** has the value "False", the IS shall perform the action indicated by table 6.

Table 5 – Level 1 state table for matching areas

Circuit Type (1)	Adjacency Usage	
	none (2)	Level 1 (3)
Level 1 only	Up (4) L1 (5)	Accept
Level 2 only	Reject (7) (Wrong system)	Down (6) (Wrong system)
Level 1 and 2	Up (4) L1 (5)	Accept
(1) The value of the Circuit Type field in the received PDU. (2) The adjacency is not in adjacencyState "Up". (3) adjacencyUsage is "Level 1" (4) The adjacency is accepted and an adjacencyStateChange (Up)" event is generated. If the Adjacency neighbourSystemType was "Unknown" (i.e. no ISH PDU has yet been received), a point-to-point IIH PDU is also transmitted. (5) The adjacencyUsage status is set to "Level 1" (6) An adjacencyStateChange (Down)" event is generated, with the specified reason, and the adjacency deleted. (7) A wrongSystemType event is generated.		

- 4) If the local system is of **iSType** "L2Intermediate-System" and the Circuit **manualL2OnlyMode** has the value "True", the IS shall perform the action indicated by table 7.
- b) If a no match is detected between any pair, the following actions shall be performed.
- 1) If the local system is of **iSType** "L1Intermediate-System" and the adjacency is not in **state** "Up", the IS shall delete the adjacency (if any) and generate an **areaMismatch** event.
 - 2) If the local system is of **iSType** "L1IntermediateSystem" and the adjacency is in **state** "Up", the IS shall delete the adjacency and generate an **adjacencyStateChange** (Down – Area Mismatch)" event .
 - 3) If the local system is of **iSType** "L2IntermediateSystem" the IS shall perform the action indicated by table 8 (irrespective of the value of **manuall2OnlyMode** for this circuit).
- c) If the action taken is "Up", as detailed in the tables referenced above, the IS shall compare the **Source ID** field of the PDU with the local **systemID**.
- 1) If the local Intermediate system has the higher Source ID, the IS shall set the Circuit **CircuitID** status to the concatenation of the local **systemID** and the Local Circuit ID (as sent in the **Local Circuit ID** field of point-to-point IIH PDUs from this Intermediate System) of this circuit.
 - 2) If the remote Intermediate system has the higher Source ID, the IS shall set the Circuit **CircuitID** status to the concatenation of the remote system's Source ID (from the **Source ID** field of the PDU), and the remote system's Local Circuit ID (from the **Local Circuit ID** field of the PDU).
 - 3) If the two source IDs are the same (i.e. the system is initialising to itself), the local **systemID** is used.
- NOTE 47 The **circuitID** status is not used to generate the Local Circuit ID to be sent in the **Local Circuit ID** field of IIH PDUs transmitted by this Intermediate system. The Local Circuit ID value is assigned once, when the circuit is created and is not subsequently changed.
- d) If the action taken is "Accept" and the neighbor System ID in the adjacency does not match the source ID field from the PDU, or the new value computed for the circuit ID is different from that in the existing adjacency, the IS shall
- 1) generate an **adjacencyStateChange** (Down) event, and
 - 2) delete the adjacency.
- e) If the action taken is "Up" or "Accept" the IS shall
- 1) copy the adjacency **areaAddressesOfNeighbour** entries from the **Area Addresses** field of the PDU,
 - 2) set the **holdingTimer** to the value of the **Holding Time** field from the PDU, and
 - 3) set the **neighbourSystemID** to the value of the **Source ID** field from the PDU.

Table 6 – Level 2 state table for matching areas

Circuit Type (1)	Adjacency Usage			
	none (2)	Level 1 (3)	Level 1 and 2 (4)	Level 2
Level 1 only	Up (6) L1 (7)	Accept	Down (8) (Wrong system)	Down (8) (Wrong system)
Level 2 only	Up (6) L2O (9)	Down (8) (Wrong system)	Down (6) (Wrong system)	Accept
Level 1 and 2	Up (6) L2 (10)	Down (8) (Wrong system)	Accept	Down (8) (Wrong system)

(1) The value of the **Circuit Type** field in the received PDU.
(2) The adjacency is not in **adjacencyState** “Up”.
(3) The adjacency is in state “Up” and the Adjacency **adjacencyUsage** is “Level 1”.
(4) The adjacency is in state “Up” and the Adjacency **adjacencyUsage** is “Level 1 and 2”.
(5) The adjacency is in state “Up” and the Adjacency **adjacencyUsage** is “Level 2”.
(6) The adjacency is accepted and an **adjacencyStateChange** (Up)” event is generated. If the Adjacency **neighbourSystemType** was “Unknown” (i.e. no ISH PDU has yet been received), a point-to-point IIH PDU is also transmitted.
(7) The **adjacencyUsage** status is set to “Level 1”.
(8) An **adjacencyStateChange** (Down)” event is generated, with the specified reason, and the adjacency deleted.
(9) The **adjacencyUsage** is set to “Level 2”.
(10) The **adjacencyUsage** is set to “Level 1 and 2”.

Table 7 – Level 2 only state table for matching areas

Circuit Type (1)	Adjacency Usage		
	none (2)	Level 1 and 2 (3)	Level 2 (4)
Level 1 only	Reject (5) (Wrong system)	Down (6) (Wrong system)	Down (6) (Wrong system)
Level 2 only	Up (7) L2O (8)	Down (6) (Wrong system)	Accept
Level 1 and 2	Up (7) L2O (8)	Down (6) (Wrong system)	Accept

(1) The value of the **Circuit Type** field in the received PDU.
(2) The adjacency is not in **adjacencyState** “Up”.
(3) The adjacency is in state “Up” and the **adjacencyUsage** is “Level 1 and 2”.
(4) The adjacency is in **adjacencyState** “Up” and the **adjacencyUsage** is “Level 2”.
(5) A **wrongSystemType** event is generated.
(6) An **adjacencyStateChange** (Down)” event is generated, with the specified reason, and the adjacency deleted.
(7) The adjacency is accepted and an **adjacencyStateChange** (Up)” event is generated. If the Adjacency **neighbourSystemType** was “Unknown” (i.e. no ISH PDU has yet been received), a point-to-point IIH PDU is also transmitted.
(8) The **adjacencyUsage** status is set to “Level 2”.

8.2.6 Monitoring point-to-point adjacencies

The IS shall keep a holding time (**adjacency holdingTimer**) for the point-to-point adjacency. The value of the **holdingTimer** shall be set to the holding time as reported in the **Holding Time** field of the Pt-Pt IIH PDU. If a neighbour is not heard from in that time, the IS shall

- purge it from the database; and
- generate an **adjacencyStateChange** (Down) event.

8.3 ISO 8208 subnetworks

8.3.1 Network layer protocols

The way in which the underlying service assumed by ISO 8473 is provided for ISO 8208 subnetworks is described in clause 8 of ISO 8473. This defines a set of Subnetwork Dependent Convergence Functions (SND CFs) that relate the service provided by specific individual ISO International Standard subnetworks to the abstract “underlying service” defined in 5.5 of ISO 8473. In particular 8.4.3 describes the Subnetwork Dependent Convergence Functions used with ISO 8208 Subnetworks.

Table 8 – Level 2 state table for non-matching areas

Circuit Type (1)	Adjacency Usage			
	none (2)	Level 1 (3)	Level 1 and 2 (4)	Level 2 (5)
Level 1 only	Reject (6) (Area Mismatch)	Down (7) (Area Mismatch)	Down (7) (Wrong system)	Down (7) (Wrong system)
Level 2 only	Up (8) L2O (9)	Down (7) (Wrong system)	Down (7) (Wrong system)	Accept
Level 1 and 2	Up (8) L2O (9)	Down (7) (Wrong system)	Down (7) (Area Mismatch)	Accept
(1) The value of the Circuit Type field in the received PDU. (2) The adjacency is not in adjacencyState “Up”. (3) The adjacency is in adjacencyState “Up” and the Adjacency adjacencyUsage is “Level 1”. (4) The adjacency is in adjacencyState “Up” and the Adjacency adjacencyUsage is “Level 1 and 2”. (5) The adjacency is in adjacencyState “Up” and the Adjacency adjacencyUsage is “Level 2”. (6) An areaMismatch event is generated. (7) An adjacencyStateChange (Down)” event is generated, with the specified reason, and the adjacency deleted. (8) The adjacency is accepted and an adjacencyStateChange (Up)” event is generated. If the Adjacency neighbourSystemType was “Unknown” (i.e. no ISH PDU has yet been received), a point-to-point ITH PDU is also transmitted. (9) The adjacencyUsage is set to “Level 2”				

8.3.2 SVC establishment

8.3.2.1 Use of ISO 8473 subnetwork dependent convergence functions

SVCs shall be established according to the procedures defined in the ISO 8208 Subnetwork Dependent Convergence Functions of ISO 8473 (this may be on system management action or on arrival of data depending on the type of circuit). The Call Request shall contain a Protocol Discriminator specifying ISO 8473 in the first octet of Call Userdata.

In the case of a *static* circuit, an SVC shall be established only upon system management action. The IS shall use **neighbourSNPAAddress** as the called SNPA address.

In the case of a DA circuit, the call establishment procedures are initiated by the arrival of traffic for the circuit.

8.3.2.2 Dynamically assigned circuits

A dynamically assigned circuit has multiple adjacencies, and can therefore establish SVCs to multiple SNPAs. There are several methods that can be used by an Intermediate system to derive the SNPA address to which a call is to be established when an NPDU is to be forwarded over an ISO 8208 subnetwork. These include the following:

- In some instances, the SNPA address to which a call is to be established can be derived from the NSAP to which an NPDU is to be forwarded.

In the case where all the NSAPs accessible over the ISO 8208 subnetwork have IDIs which are their SNPA addresses, the correct SNPA can be derived by extracting the IDI, using the “extractIDI” mapping type described in 8.1.

Other scenarios may also permit the extraction of the SNPA by examining other parts of the NSAP address. In these cases the “extractDSP” mapping type may be used as described in 8.1.

Examples of the above methods are illustrated in table 3.

- In other cases, such as when the IDI refers to some other SNPA address which is suboptimally connected to the target NSAP (or even not connected at all), or when the IDP does not contain an X.121 address at all (e.g. the ISO DCC address plan), a method not relying upon information in the destination NSAP address must be used.

If it is feasible for the IS to maintain the correspondence between an address prefix and an SNPA (via the Reachable Address managed object) then the “explicit” mapping type may be used as described in 8.1. This may not always be desirable because of the need to administer this information individually in each affected Intermediate system.

If a SNARE is available on the subnetwork, then the IS may invoke the appropriate SNARE functions to obtain the desired SNPA address from the NSAP address in the NPDU to be forwarded.

This is achieved, as described in 8.1, by additional information contained in the **reachableAddress** managed object. The address extraction algorithm may be specified to extract the IDI or DSP portion where the desired portion of the destination NSAP address is the required X.121 address. An example of a set of Reachable Addresses is shown in table 3.

NOTE 48 If a DA circuit is defined with a reachable address prefix which includes the addresses reachable over a STATIC circuit, the cost(s) for the DA circuit must be greater than those of the STATIC circuit. If this is not the case, the DA circuit may be used to establish a call to the remote SNPA supporting the STATIC circuit, which would then (wrongly) assume it was the STATIC circuit.

8.3.2.3 Initiating calls (level 2 Intermediate systems)

When an NPDU is to be forwarded on a dynamically assigned circuit, for destination NSAP address D, the IS shall

- a) Calculate D's subnetwork address, either as explicitly stated in the reachable address prefix, or as extracted from the destination NSAP address.
 - 1) If this system is an ES and there is an entry in the RedirectCache or ReversePathCache for D, use the subnetwork address in the cache entry.
 - 2) If this system is an ES or Level 2 Intermediate system, and the address matches one of the listed reachable address prefixes (including "*", if present), the subnetwork address is that specified according to the **mappingType** attribute (either "explicit", indicating that the set of addresses in the **sNPAAAddresses** attribute of that Reachable Address are to be used, or "Algorithm", indicating that it is to be extracted from the destination NSAP address using the specified algorithm). If multiple SNPA addresses are specified, and there is already an adjacency up to one of those SNPA addresses, then choose that subnetwork address, otherwise choose the subnetwork address with the oldest timestamp as described in 8.3.2.4.
 - 3) If the address does not match one of the listed reachable address prefixes (and there is no "*" entry), invoke the ISO 8473 Discard PDU function.
- b) Scan the adjacencies for one already open to D's subnetwork address (i.e. **reserveTimer** has not yet expired). If one is found, transmit the NPDU on that adjacency.
- c) If no adjacency has a call established to the required subnetwork address, but there is a free adjacency, attempt to establish the call using that subnetwork address.
- d) If there is no free adjacency invoke the ISO 8473 Discard PDU function.

NOTE 49 Where possible, when an adjacency is reserved (when an SVC has been cleared as a result of the **idleTimer** expiring, but the **reserveTimer** has not yet expired), resources within the subnetwork service provider should be reserved, in order to minimise the probability that the adjacency will not be able to initiate a call when required.

8.3.2.4 Call attempt failures

The Reachable Address managed objects may contain a set of SNPA addresses, each of which has an associated timestamp. The timestamps shall be initialised to "infinitely old".

Some of the SNPAs in this set may be unreachable. If a call attempt fails to one of the SNPA addresses listed, the IS shall mark that entry in the list with the time of the latest failed attempt. When an SNPA address is to be chosen from the list, the IS shall choose the one with the oldest timestamp, unless the oldest timestamp is more recent than **recallTimer**. If the oldest timestamp is more recent than **recallTimer**, all SNPAs in the set shall be assumed temporarily unreachable and no call attempt is made. The IS shall instead invoke the ISO 8473 Discard PDU function.

When attempting to establish a connection to a single specific subnetwork address (not through one of a set of SNPA addresses), if a call attempt to a particular SNPA address, A, fails for any reason, the IS shall invoke the ISO 8473 Discard PDU function. Additionally the adjacency on which the call attempt was placed shall be placed in "Failed" state, and the recall timer set. Until it expires, the IS shall not attempt call establishment for future NPDUs to be forwarded over subnetwork address A, but instead the IS shall invoke the ISO 8473 Discard PDU function.

When the recall timer expires, the IS shall free the adjacency for calls to a different destination or retry attempts to subnetwork address A.

NOTE 50 If an implementation can store the knowledge of SNPA addresses that have failed along with the time since the attempt was made in a location other than the adjacency on which the call was attempted, then that adjacency can be used for other calls.

8.3.3 Reverse path forwarding on DA circuits

Where a subdomain is attached to a Connection-oriented subnetwork by two or more SNPAs, the destination NSAP addresses within the subdomain may be chosen to be constructed from the address of one of the points of attachment. (It need not be. The whole subdomain could be multi-homed by using both SNPA addresses, or some other IDP could be chosen; e.g. ISO DCC.) Traffic to the subdomain from some other SNPA will cause a call to be established to the SNPA corresponding to the destination NSAP

address in the subdomain. Traffic *from* the subdomain may use either of the SNPAs depending on the routing decisions made by the subdomain. This is illustrated in figure 6.

The subdomain is attached to the connection-oriented subnetwork via SNPAs *A* and *B*. The addresses on the subdomain are constructed using the SNPA address of *B* as the IDI. If traffic for *C.z* is sent from *B.x*, a call will be established from *A* to *C*. The reverse traffic from *C.z* to *B.x* will cause another call to be established from *C* to *B*. Thus two SVCs have been established where only one is required.

This problem is prevented by the local system retaining a cache (known as the ReversePathCache) of NSAP addresses from which traffic has been received over each adjacency. When it has traffic to forward over the connection-oriented subnetwork, the IS shall first check to see if the destination NSAP is in the cache of any of its adjacencies, and if so forwards the traffic over that adjacency. An NSAP shall only be added to the cache when the remote SNPA address of the adjacency over which it is received differs from the SNPA address to be called which would be generated by checking against the Circuit Reachable Addresses managed objects. If the cache is full, the IS shall overwrite the least recently used entry. The ReversePathCache, if implemented, shall have a size of at least one entry. The IS shall purge the cache when the adjacency is taken down (i.e. when the reserveTimer expires).




Figure 6 - Example of reverse path forwarding

8.3.4 Use of ISO 9542 on ISO 8208 subnetworks

Static and DA circuits are equivalent to point-to-point links, and as such permit the operation of ISO 9542 as described for point-to-point links in 8.2.

For DA circuits, it is impractical to use ISO 9542 to obtain configuration information, such as the location of Intermediate systems, since this would require calls to be established to all possible SNPA addresses.

The IS shall not send ISO 9542 ISH PDUs on a DA circuit. The IS shall take no action on receipt of an ESH PDU or ISH PDU, and the circuit shall complete initialisation without waiting for their arrival.

The IS shall not send Point-to-point IIH PDU on DA circuits.

The IS shall ignore receipt of point-to-point IIH PDUs on DA circuits.

8.3.5 Interactions with the update process

A dynamically assigned circuit contains a list of <reachable address prefix, cost, SNPA address> tuples. Also, each dynamically assigned circuit has a specified call establishment cost measured by $\text{callEstablishmentMetric}_k$ (where k indexes the four defined metrics). The call establishment cost is always an internal metric, and is therefore directly comparable with the reachable address metric only if the reachable address metric is also internal.

When the circuit is enabled, the Subnetwork Dependent functions in an Intermediate system shall report (to the Update Process) “adjacency cost change” events for all address prefixes in the circuit reachable address managed object, together with the reachable address metric $k + \Delta_k$ increment. If reachable address metric_k is internal, then $\Delta_k = \text{callEstablishmentMetric}_k$. If reachable address metric_k is external, then $\Delta_k = 0$.

This causes this information to be included in subsequently generated LSPs as described in 7.3.3.2. Routeing PDUs (LSPs and Sequence number PDUs) shall not be sent on dynamically assigned circuits.

NOTE 51 In the following sub-clauses, it is assumed that the Reachable Addresses referenced are only those which have been enabled (i.e. that have state “On”), and whose parent circuit is also in state “On”.

8.3.5.1 Adjacency creation

After an SVC to SNPA address D is successfully established and a new adjacency created for it (whether it was initiated by the local or the remote system), if $\text{callEstablishment-Metric}_k\text{Increment}$ is greater than 0, the IS shall scan the circuit Reachable Address managed objects for all **addressPrefixes** listed with D as (one of) the **SNPAAddress(es)**.

For Reachable Addresses with **mappingType** “Algorithm”, the IS shall construct an *implied address prefix*¹⁾ from the actual remote SNPA address D and the address extraction algorithm. The IS shall generate an Adjacency cost change event for each such address prefix (both actual and implied) with the Reachable Address metric_k (without the added $\text{callEstablishment-Metric}_k\text{Increment}$). This causes information that those address prefixes are reachable with the lower cost to be included in subsequently generated LSPs. The effect of this is to encourage the use of already established SVCs where possible.

8.3.5.2 Adjacency deletion

When the adjacency with **SNPAAddress** D is freed (reserveTimer has expired, or the adjacency is deleted by System Management action) then if $\text{callEstablishment-Metric}_k\text{Increment}$ is greater than 0, the IS shall scan the Circuit Reachable Address managed objects for all those with **mappingType** **explicit** and (one of) their **SNPAAddresses** equal to D . The IS shall generate “Adjacency cost change” events to the Update Process for all such address prefixes with the Reachable Address metric_k + Δ_k increment (where Δ_k is the same as defined above). For Reachable Addresses with an algorithmic extraction **mappingType** for which it is possible to construct an implied address prefix as above, the IS shall generate an **adjacencyStateChange** event for that implied prefix.

A cost change event shall only be generated when the count of the number of subnetwork addresses which have an established SVC changes between 1 and 0.

8.3.5.3 Circuit call establishment increment change

On a dynamically assigned circuit, when system management changes the Circuit $\text{callEstablishmentMetric}_{kl}\text{Increment}$ for that circuit, the IS shall generate “adjacency cost change events” for all address prefixes affected by the change (i.e. those for which calls are not currently “established”).

The IS shall scan all the Reachable Address managed objects of that Circuit. If the Reachable Address has an algorithmic

extraction **mappingType**, the IS shall generate an “adjacency cost change” event for that **adjacencyId** with the Reachable Address metric_k + the new value of Δ_k . If (based on the new value of $\text{callEstablishmentMetric}_k\text{Increment}$) the Reachable Address has **mappingType** “explicit”, the IS shall scan all the adjacencies of the circuit for an adjacency with **SNPAAddress** equal to (one of) the **SNPAAddresses** of that Reachable Address. If no such adjacency is found the IS shall generate an “adjacency cost change” event for that **adjacencyId** with the reachable address metric_k + the new value of Δ_k (based on the new value of $\text{callEstablishmentMetric}_k\text{Increment}$).

8.3.5.4 Reachable address cost change

When the metric_k attribute of a **reachableAddress** in **operationalState** “Enabled” is changed by system management, the IS shall generate cost change events to the Update process to reflect this change.

¹⁾ i.e. some address prefix which matches the **addressPrefix** of the Reachable Address, and which would generate the SNPA Address D when the extraction algorithm is applied.

If the **reachableAddress** has **mappingType** “explicit”, the IS shall scan all the adjacencies of the circuit for an adjacency with **sNPAAddress** equal to (one of) the **sNPAAddresses** of that reachable address. If one or more such adjacencies are found, the IS shall generate an **adjacencyCostChange** event for that **reachableAddressId** with the new reachable address metric_k. If no such adjacency is found the IS shall generate an “adjacency cost change” event for that **reachableAddressId** with the new reachable address metric_k.

If the **reachableAddress** has an algorithmic extraction **mappingType**, the IS shall generate an **adjacencyCostChange** event for that **name** with the new reachable address metric_k + Delta_k (based on the new value of **callEstablishment-Metric_kIncrement**). In addition, for all adjacencies of the circuit with an **sNPAAddress** for which an implied address prefix can be generated for this reachable address, the IS shall generate an **adjacencyCostChange** event for that implied address prefix and the new reachable address metric_k.

8.3.5.5 Disabling a reachable address

When a **reachableAddress** managed object is disabled via management action, the IS shall generate an **adjacency-Down** event to the Update process for the **adjacencyId** of that reachable address and also for any implied prefixes associated with that reachable address.

8.3.5.6 Enabling a reachable address

When a **reachableAddress** is enabled via system management action, the IS shall generate **adjacencyCostChange** events as described for reachable address cost changes in 8.3.5.4 above.

8.4 Broadcast subnetworks

8.4.1 Enabling of broadcast circuits

When the broadcast circuit is enabled on an Intermediate system the IS shall perform the following actions.

- a) Commence sending IIH PDUs with the LAN ID field set to the concatenation of its own **systemID** and its locally assigned one octet Local Circuit ID.
- b) Solicit the End system configuration as described in 8.4.6
- c) Start listening for ISO 9542 ESH PDUs and acquire adjacencies as appropriate. Do not run the Designated Intermediate System election process.
- d) After waiting **iSISHelloTimer** × 2 seconds, run the Level 1 and or the level 2 designated intermediate system election process depending on the Intermediate system type.

8.4.2 Broadcast subnetwork IIH PDUs

All Intermediate systems on broadcast circuits (both Level 1 and Level 2) shall transmit LAN IIH PDUs as described in 8.4.4. Level 1 Intermediate systems shall transmit only Level 1 LAN IIH PDUs. Level 2 Intermediate Systems on circuits with **manualL2OnlyMode** set to the value “True”, shall transmit only Level 2 LAN IIH PDUs. Level 2 Intermediate systems on circuits with **manualL2OnlyMode** set to the value “False”, shall transmit both.

Level *n* LAN IIH PDUs contain the transmitting Intermediate system’s ID, holding timer, Level *n* Priority and **manualAreaAddresses**, plus a list containing the **lanAddresses** of all the adjacencies of neighbourSystemType “Ln Intermediate System” (in **adjacencyState** “Initialising” or “Up”) on this circuit.

The Circuit Type field shall be set according to Table 4. In a level 1 IIH PDU the Circuit Type shall be either 1 or 3. In a level 2 IIH PDU the Circuit Type shall be either 2 or 3.

LAN IIH PDUs shall be padded (with trailing PAD option fields containing arbitrary valued octets) so that the SNSDU containing the IIH PDU has a length of at least *maxsize* - 1 octets¹⁾ where *maxsize* for Level 1 IIH PDUs is the maximum of

- **dataLinkBlockSize**
- **originatingL1LSPBufferSize**

¹⁾ The minimum length of PAD which may be added is 2 octets, since that is the size of the option header. Where possible the PDU should be padded to *maxsize*, but if the PDU length is *maxsize*-1 octets no padding is possible (or required).

and for Level 2 IIH PDUs is the maximum of

- dataLinkBlockSize
- originatingL2LSPBufferSize

This is done to ensure that an adjacency will only be formed between systems which are capable of exchanging PDUs of length up to *maxsize* octets. In the absence of this check, it would be possible for an adjacency to exist with a lower maximum block size, with the result that some LSPs and SNPs (i.e. those longer than this maximum, but less than *maxsize*) would not be exchanged.

NOTE 52 An example of a topology where this could occur is one where an extended LAN is constructed from LAN segments with different maximum block sizes. If, as a result of mis-configuration or some dynamic reconfiguration, a path exists between two Intermediate systems on separate LAN segments having a large maximum block size, which involves transit of a LAN segment with a smaller maximum block size, loss of larger PDUs will occur if the Intermediate systems continue to use the larger maximum block size. It is better to refuse to bring up the adjacency in these circumstances.

Level 1 Intermediate systems shall transmit Level 1 LAN IIH PDUs to the multi-destination address AllL1ISs, and also listen on that address. They shall also listen for ESH PDUs on the multi-destination address AllIntermediateSystems. The list of neighbour Intermediate systems shall contain only Level 1 Intermediate Systems within the same area. (i.e. Adjacencies of neighbourSystemType "L1 Intermediate System".)

Level 2 Only Intermediate systems (i.e. Level 2 Intermediate systems which have the Circuit with an associated linkage manualL2OnlyMode characteristic set to the value "True") shall transmit Level 2 LAN IIH PDUs to the multi-destination address AllL2ISs, and also listen on that address. The list of neighbour Intermediate systems shall contain only Level 2 Intermediate systems. (i.e. adjacencies of neighbourSystemType "L2 Intermediate System".)

Level 2 Intermediate systems (with manualL2OnlyMode "False") shall perform both of the above actions. Separate Level 1 and Level 2 LAN IIH PDUs shall be sent to the multi-destination addresses AllL1ISs and AllL2ISs describing the neighbour Intermediate systems for Level 1 and Level 2 respectively. Separate adjacencies shall be created by the receipt of Level 1 and Level 2 LAN IIH PDUs.

8.4.2.1 IIH PDU acceptance tests

On receipt of a Broadcast IIH PDU, perform the following PDU acceptance tests:

- a) If the IIH PDU was received over a circuit whose externalDomain attribute is "True", the IS shall discard the PDU.
- b) If the ID Length field of the PDU is not equal to the value of the IS's routingDomainIDLength, the PDU shall be discarded and an idFieldLengthMismatch event generated.
- c) If the value of circuitTransmitPassword or the set of circuitReceivePasswords for this circuit is non-null, then perform the following tests:
 - 1) If the PDU does not contain the Authentication Information field then the PDU shall be discarded and an authenticationFailure event generated.
 - 2) If the PDU contains the Authentication Information field, but the Authentication Type is not equal to "Password", then
 - i. If the IS implements the authentication procedure indicated by the Authentication Type whether the IS accepts or ignores the PDU is outside the scope of this International Standard.
 - ii. If the IS does not implement the authentication procedure indicated by the Authentication Type then the IS shall ignore the PDU and generate an authenticationFailure event."
 - 3) Otherwise, the IS shall compare the password in the received PDU with the passwords in the set of circuitReceivePasswords for the circuit on which the PDU was received. If the value in the PDU matches any of these passwords, the IS shall accept the PDU for further processing. If the value in the PDU does not match any of the circuitReceivePasswords, then the IS shall ignore the PDU and generate an authenticationFailure event.

8.4.2.2 Receipt of level 1 IIH PDUs

On receipt of a Level 1 LAN IIH PDU on the multi-destination address AllL1ISs, the IS shall perform the following tests:

- a) Compare each of the area addresses, from the **Area Addresses** field of the received IIH PDU with the set of area addresses in the **manualAreaAddresses** attribute. If a match is not found between any pair (i.e. the local and remote system have no area address in common), the IS shall reject the adjacency and generate an **areaMismatch** event.
- b) if the **Maximum Area Addresses** field of the PDU is not equal to the value of the IS's **maximumAreaAddresses** then the PDU shall be discarded and a **maximumArea-AddressesMismatch** event generated, unless the IS only implements a value of three for **maximumAreaAddresses**, in which case this check may be omitted. If the above tests succeed, the IS shall accept the adjacency and set the Adjacency **neighbourSystemType** to "L1 Intermediate System".

8.4.2.3 Receipt of Level 2 IIH PDUs

On receipt of a Level 2 LAN IIH PDU on the multi-destination address **AllL2ISs**, the IS shall accept the adjacency, and set the Adjacency **neighbourSystemType** to "L2 Intermediate System".

8.4.2.4 Existing adjacencies

When a Level n LAN IIH PDU (Level 1 or Level 2) is received from an Intermediate system for which there is already an adjacency with

- a) the adjacency **neighbourSNPAAAddress** equal to the MAC Source address of the PDU, **and**
- b) the Adjacency **neighbourSystemID** equal to the Source ID field from the PDU **and**
- c) the **neighbourSystemType** equal to " L_n Intermediate System", the IS shall update the **holdingTimer**, **priorityOfNeighbour** and **areaAddressesOfNeighbour** according to the values in the PDU.

8.4.2.5 New adjacencies

8.4.2.5.1 When

- a) Level n LAN IIH PDU (Level 1 or Level 2) is received (from Intermediate system R), and
- b) there is no adjacency for which the adjacency **neighbourSNPAAAddress** is equal to the MAC Source address of the PDU; and

the IS shall create a new adjacency. However, if there is insufficient space in the adjacency database, to permit the creation of a new adjacency the IS shall instead perform the actions described in 8.4.3.

The IS shall

- a) set **neighbourSystemType** to " L_n Intermediate System" (where n is the level of the IIH PDU),
- b) set the **holdingTimer**, **priorityOfNeighbour**, **neighbour-SystemID** and **areaAddressesOfNeighbour** according to the values in the PDU., and
- c) set the **neighbourSNPAAAddress** according to the MAC source address of the PDU.

The IS shall set the **adjacencyState** of the adjacency to "initialising", until it is known that the communication between this system and the source of the PDU (R) is two-way. However R shall be included in future Level n LAN IIH PDUs transmitted by this system.

When R reports the local system's SNPA address in its Level n LAN IIH PDUs, the IS shall

- d) set the adjacency's **adjacencyState** to "Up", and
- e) generate an **adjacencyStateChange (Up)** event.

8.4.2.5.2 The IS shall keep a separate holding time (adjacency **holdingTimer**) for each " L_n Intermediate System" adjacency. The value of **holdingTimer** shall be set to the holding time as reported in the **Holding Time** field of the Level n LAN IIH PDUs. If a neighbour is not heard from in that time, the IS shall

- a) purge it from the database; and
- b) generate an **adjacencyStateChange (Down)** event.

8.4.2.5.3 If a Level n LAN IIH PDU is received from neighbour N , and this system's LANAddress is no longer in N 's IIH PDU, the IS shall

- a) set the adjacency's adjacencyState to "initialising", and
- b) generate an adjacencyStateChange (Down) event.

8.4.3 Insufficient space in adjacency database

If an IS needs to create a new Intermediate system adjacency, but there is insufficient space in the adjacency database, the adjacency of neighbourSystemType " L_n Intermediate System" with lowest $l \times$ intermediateSystemPriority (or if more than one adjacency has the lowest priority, the adjacency with the lowest SNPAAddress, from among those with the lowest priority) shall be purged from the database. If the new adjacency would have the lowest priority, it shall be ignored, and a rejectedAdjacency event generated.

If an old adjacency must be purged, the IS shall generate an adjacencyStateChange (Down) event for that adjacency. After the Subnetwork Independent Functions issue an "adjacency down complete", the IS may create a new adjacency.

8.4.4 Transmission of LAN IIH PDUs

A Level 1 IS shall transmit a Level 1 LAN IIH PDU immediately when any circuit whose externalDomain attribute is "False" has been enabled. A Level 2 Intermediate System shall transmit a Level 2 LAN IIH PDU. A Level 2 Intermediate System shall also transmit a Level 1 LAN IIH PDU unless the circuit is marked as manualL2OnlyMode "True".

The IS shall also transmit a LAN IIH PDU when at least 1 second has elapsed since the last transmission of a LAN IIH PDU of the same type on this circuit by this system and:

- a) iSISHelloTimer seconds have elapsed¹⁾ since the last periodic LAN IIH PDU transmission.

The Holding Time is set to $iSISHoldingMultiplier \times iSISHelloTimer$. For a Designated Intermediate System the value of $dRISISHelloTimer$ ²⁾ is used instead of $iSISHelloTimer$. The Holding Time for this PDU shall therefore be set to $iSISHoldingMultiplier \times dRISIS-HelloTimer$ seconds. This permits failing Designated In-termediate Systems to be detected more rapidly,

or

- b) the contents of the next IIH PDU to be transmitted would differ from the contents of the previous IIH PDU transmitted by this system, or
- c) this system has determined that it is to become or resign as LAN Designated Intermediate System for that level.

To minimise the possibility of the IIH PDU transmissions of all Intermediate systems on the LAN becoming synchronised, the hello timer shall only be reset when a IIH PDU is transmitted as a result of timer expiration, or on becoming or resigning as Designated Intermediate System.

Where an Intermediate system is transmitting both Level 1 and Level 2 LAN IIH PDUs, it shall maintain a separate timer (separately jittered) for the transmission of the Level 1 and Level 2 IIH PDUs. This avoids correlation between the Level 1 and Level 2 IIH PDUs and allows the reception buffer requirements to be minimised.

If the value of the circuitTransmitPassword for the circuit is non-null, then the IS shall include the Authentication Information field in the transmitted IIH PDU, indicating an Authentication Type of "Password" and containing the circuitTransmitPassword as the authentication value.

8.4.5 LAN designated intermediate systems

A LAN Designated Intermediate System is the highest priority Intermediate system in a particular set on the LAN, with numerically highest MAC source SNPAAddress breaking ties. (See 7.1.8 for how to compare LAN addresses.)

¹⁾ Jitter is applied as described in 10.1.

²⁾ In this case jitter is not applied, since it would result in intervals of less than one second.

There are, in general, two LAN Designated Intermediate Systems on each LAN, namely the LAN Level 1 Designated Intermediate System and the LAN Level 2 Designated Intermediate System. They are elected as follows.

- a) Level 1 Intermediate systems elect the LAN Level 1 Designated Intermediate System.
- b) Level 2 Only Intermediate systems (i.e. Level 2 Intermediate Systems which have the Circuit `manualL2OnlyMode` characteristic set to the value “True”) elect the LAN Level 2 Designated Intermediate System.
- c) Level 2 Intermediate systems (with `manualL2OnlyMode` “False”) elect both the LAN Level 1 and LAN Level 2 Designated Intermediate Systems.

The set of Intermediate systems to be considered includes the local Intermediate system, together with all Intermediate systems of the appropriate type as follows.

- d) For the LAN Level 1 Designated Intermediate System, it is the set of Intermediate systems from which LAN Level 1 IIH PDUs are received and to which Level 1 adjacencies exist in `adjacencyState` “Up”. When the local system either becomes or resigns as LAN Level 1 Designated Intermediate System, the IS shall generate a `LANLevel1DesignatedIntermediateSystemChange` event. In addition, when it becomes LAN Level 1 Designated Intermediate System, it shall perform the following actions.
 - 1) Generate and transmit Level 1 pseudonode LSPs using the existing End system configuration.
 - 2) Purge the Level 1 pseudonode LSPs issued by the previous LAN Level 1 Designated Intermediate System (if any) as described in 7.2.3.
 - 3) Solicit the new End system configuration as described in 8.4.6.
- e) For the LAN Level 2 Designated Intermediate System, it is the set of Intermediate systems from which LAN Level 2 IIH PDUs are received and to which Level 2 adjacencies exist in `adjacencyState` “Up”. When the local system either becomes or resigns as LAN Level 2 Designated Intermediate System, the IS shall generate a `LANLevel2DesignatedIntermediateSystemChange` event. In addition, when it becomes LAN Level 2 Designated Intermediate System, it shall perform the following actions.
 - 1) Generate and transmit a Level 2 pseudonode LSP.
 - 2) Purge the Level 2 pseudonode LSPs issued by the previous LAN Level 2 Designated Intermediate System (if any) as described in 7.2.3.

When an Intermediate system resigns as LAN Level 1 or Level 2 Designated Intermediate System it shall perform the actions on Link State PDUs described in 7.2.3.

The IS shall run the Level 1 and/or the Level 2 Designated Intermediate System election process (depending on the Intermediate system type) whenever an IIH PDU is received or transmitted as described in 8.4.4. (For these purposes, the transmission of the system's own IIH PDU is equivalent to receiving it). If there has been no change to the information on which the election is performed since the last time it was run, the previous result can be assumed. The relevant information is

- f) the set of Intermediate system adjacency states;
- g) the set of Intermediate System priorities (including this system's); and
- h) the existence (or otherwise) of at least one “Up” End system (not including manual adjacencies) or Intermediate system adjacency on the circuit.

An Intermediate system shall not declare itself to be a LAN Designated Intermediate system of any type until it has at least one “Up” End system (not including manual adjacencies) or Intermediate system adjacency on the circuit. (This prevents an Intermediate System which has a defective receiver and is incapable of receiving any PDUs from erroneously electing itself LAN Designated Intermediate System.)

The LAN ID field in the LAN IIH PDUs transmitted by this system shall be set to the value of the LAN ID field reported in the LAN IIH PDU (for the appropriate level) received from the system which this system considers to be the Designated Intermediate System. This value shall also be passed to the Update Process, as the pseudonode ID, to enable Link State PDUs to be issued for this system claiming connectivity to the pseudonode.

If this system, as a result of the Designated Intermediate System election process, considers itself to be designated Intermediate System, the LAN ID field shall be set to the concatenation of this system's own system ID and the locally assigned one octet Local Circuit ID.

8.4.6 Soliciting the ES configuration

When there is a change in the topology or configuration of the LAN (for example the partitioning of a LAN into two segments by the failure of a repeater or bridge), it is desirable for the (new) Designated Intermediate System to acquire the new End system configuration of the LAN as quickly as possible in order that it may generate Link State PDUs which accurately reflect the actual configuration. This is achieved as follows.

When the circuit is enabled, or the Intermediate system detects a change in the set of Intermediate systems on the LAN, or a change in the Designated Intermediate System ID, the IS shall initiate a poll of the ES configuration by performing the following actions.

- a) Delay a random interval between 0 and `iSISHelloTimer`. (This is to avoid synchronisation with other Intermediate systems which have detected the change.)
- b) If (and only if) an Intermediate System had been removed from the set of Intermediate systems on the LAN, reset the `entryRemainingTime` field in the `neighbour-SystemIDs` adjacency database record of all adjacencies on this circuit to the value

$$(\text{iSISHelloTimer} + \text{pollESHHelloRate}) \times \text{iSISHoldingMultiplier}$$
or the existing value whichever is the lower. (This causes any End systems which are no longer present on the LAN to be rapidly timed out, but not before they have a chance to respond to the poll.)
- c) Transmit `iSISHoldingMultiplier` ISH PDUs for each NET possessed by the Intermediate system with a Suggested ES Configuration Timer value of `pollESHHelloRate` at an interval between them of `iSISHelloTimer` and a holding time of `iSConfigurationTimer` $\times$ `iSISHoldingMultiplier`.
- d) Resume sending ISH PDUs at intervals of `iSConfigurationTimer` with a Suggested ES Configuration Timer value of `defaultESHHelloTimer`.

8.4.7 Receipt of ESH PDUs — database of end systems

An IS shall enter an End system into the adjacency database when an ESH PDU is received from a new data link address. If an ESH PDU is received with the same data link address as a current adjacency, but with a different NSAP address, the new address shall be added to the adjacency, with a separate timer. A single ESH PDU may contain more than one NSAP address. When a new data link address or NSAP address is added to the adjacency database, the IS shall generate an `adjacencyStateChange` (Up) event on that adjacency.

The IS shall set a timer for the value of the `Holding Time` field in the received ESH PDU. If another ESH PDU is not received from the ES before that timer expires, the ES shall be purged from the database, provided that the Subnetwork Independent Functions associated with initialising the adjacency have been completed. Otherwise the IS shall clear the adjacency as soon as those functions are completed.

When the adjacency is cleared, the Subnetwork Independent Functions shall be informed of an `adjacencyStateChange` (Down) event, and the adjacency can be re-used after the Subnetwork Independent Functions associated with bringing down the adjacency have been completed.

8.4.8 Broadcast subnetwork constants

The Intradomain IS-IS protocol exploits multicast capabilities for all IS-IS protocol exchanges on broadcast subnetworks. To ensure interoperability all systems on a given broadcast subnetwork must use the same multi-destination address bindings. For ISO 8802 subnetworks that supports 48 bit MAC addresses, 48 bit MAC addressing and the multi-destination address bindings in table 9 below shall be used.

Table 9 – Architectural constants for use with ISO 8802 subnetworks

Multi-destination Address	Binding	Description
AllL1ISs	01-80-C2-00-00-14	The multi-destination address “All L1 Intermediate Systems”
AllL2ISs	01-80-C2-00-00-15	The multi-destination address “All L2 Intermediate Systems”
AllIntermediateSystems	09-00-2B-00-00-05	The multi-destination address “All Intermediate Systems” used by ISO 9542
AllEndSystems	09-00-2B-00-00-04	The multi-destination address “All End Systems” used by ISO 9542

9 Structure and encoding of PDUs

This clause describes the PDU formats of the Intra-Domain IS-IS routing protocol.

9.1 General encoding rules

Octets in a PDU are numbered starting from 1, in increasing order. Bits in a octet are numbered from 1 to 8, where bit 1 is the least significant bit and is pictured on the right. When consecutive octets are used to represent a number, the lower octet number has the most significant value.

Fields marked **Reserved** (or simply R) are transmitted as zero, and ignored on receipt, unless otherwise noted.

Values are given in decimal. All numeric fields are unsigned integers, unless otherwise noted.

9.2 Encoding of network layer addresses

Network Layer addresses (NSAP addresses, NETs, area addresses and Address Prefixes) are encoded in PDUs according to the preferred binary encoding specified in ISO

8348; the entire address, taken as a whole is represented explicitly as a string of binary octets. This string is conveyed in its entirety in the address fields of the PDUs. The rules governing the generation of the preferred binary encoding are described in ISO 8348. The address so generated is encoded with the most significant octet (i.e. the AFI) of the address being the first octet transmitted, and the more significant semi-octet of each pair of semi-octets in the address is encoded in the more significant semi-octet of each octet (i.e. in the high order 4 bits). Thus the address /371234 is encoded as shown in figure 7.

		No. of Octets
3	7	1
1	2	1
3	4	1

Figure 7 - Address encoding example

9.3 Encoding of SNPA addresses

SNPA addresses (e.g. LAN Address) shall be encoded according to the rules specified for the particular type of subnetwork being used.

In the case of an ISO/IEC 8802 (except for ISO/IEC 8802-6 DQDB) or ISO/IEC 9314 subnetwork, the SNPA address is the 48-bit MAC address encoded as a sequence of six octets according to the “hexadecimal representation” of MAC addresses specified in ISO/IEC 10039. The SNPA address of an ISO/IEC 8802-6 DQDB shall be encoded as a sequence of eight octets containing 4-bit address type subfield, variable length padding subfield and variable length MAC service access point address subfield according to the “hexadecimal representation”.

NOTE 53 In this encoding the first bit of the binary representation of the SNPA address is the least significant bit of the first octet in the encoded sequence.

9.4 PDU types

The types of PDUs are:

- Level 1 LAN IS to IS Hello PDU
- Level 2 LAN IS to IS Hello PDU
- Point-to-Point IS to IS Hello PDU
- Level 1 Link State PDU
- Level 2 Link State PDU
- Level 1 Complete Sequence Numbers PDU
- Level 2 Complete Sequence Numbers PDU
- Level 1 Partial Sequence Numbers PDU
- Level 2 Partial Sequence Numbers PDU
-

These are described in the following subclauses.

9.5 Level 1 LAN IS to IS hello PDU

This PDU is multicast by Intermediate systems on broadcast circuits to the multi-destination address AILL1ISs. The purpose of this PDU is for Intermediate systems on broadcast circuits to discover the identity of other Level 1 Intermediate systems on that circuit. Trailing Pad option fields are inserted to make PDU Length equal to at least $maxsize - 1$ where $maxsize$ is the maximum of

- dataLinkBlocksize
 - originatingL1LSPBufferSize
- (see 8.4.2).

				No. of Octets
Intradomain Routing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
Reserved/Circuit Type				1
Source ID				ID Length
Holding Time				2
PDU Length				2
R	Priority			1
LAN ID				ID Length + 1
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routeing Protocol Discriminator — architectural constant (see table 2)
- Length Indicator – Length of the fixed header in octets
- Version/Protocol ID Extension – 1

- **ID Length** — Length of the ID field of NSAP addresses and NETs used in this routing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- **PDU Type** (bits 1 through 5) – 15.

NOTE 54 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt

- **Version** – 1
- **Reserved** — transmitted as zero, ignored on receipt
- **Maximum Area Addresses** — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumAreaAddresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.
- **Reserved/Circuit Type** – Most significant 6 bits reserved (Transmitted as zero, ignored on receipt). Low order bits (bits 1 and 2) indicate:
 - 0 – reserved value (if specified the entire PDU shall be ignored)
 - 1 — Level 1 only
 - 2 — Level 2 only (sender is Level 2 Intermediate system with **manualL2OnlyMode** set “True” for this circuit, and will use this link only for Level 2 traffic)
 - 3 – both Level 1 and Level 2 (sender is Level 2 Intermediate system, and will use this link both for Level 1 and Level 2 traffic)

NOTE 55 In a LAN Level 1 IIH PDU the **Circuit Type** shall be either 1 or 3.

- **Source ID** – the system ID of transmitting Intermediate system
- **Holding Time** – Holding Timer to be used for this Intermediate system
- **PDU Length** – Entire length of this PDU, in octets, including header
- **Reserved/Priority** – Bit 8 reserved (Transmitted as zero, ignored on receipt). Bits 1 through 7 – priority for being LAN Level 1 Designated Intermediate System. Higher number has higher priority for being LAN Level 1 Designated Intermediate System. Unsigned integer.
- **LAN ID** – a field composed the system ID (1–8 octets) of the LAN Level 1 Designated Intermediate System, plus a low order octet assigned by LAN Level 1 Designated Intermediate System. Copied from LAN Level 1 Designated Intermediate System’s IIH PDU.
- **VARIABLE LENGTH FIELDS** – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received PDU that are not recognised shall be ignored.

Currently defined codes are:

- **Area Addresses** – the set of `manualAreaAddresses` of this Intermediate System.

x CODE – 1
 x LENGTH – total length of the value field.
 x VALUE –

	No. of Octets
Address Length	1
Area Address	Address Length
Address Length	1
Area Address	Address Length

- **Address Length** – Length of **Area Address** in octets.
- **Area Address** – Area address.
- **Intermediate System Neighbours** – The set of Intermediate systems on this LAN to which adjacencies of `neighborSystemType` “L1 Intermediate System” exist in state “Up” or “Initialising” (i.e. those from which Level 1 IIH PDUs have been heard). This is permitted to appear more than once.

Two types of codes are defined:

- a) Intermediate System Neighbors with 6 octet MAC address
- x CODE – 6
 x LENGTH – total length of the value field in octets
 x VALUE –

	No. of Octets
LAN Address	6
LAN Address	6

- **LAN Address (6 octet length)** – 6 octet MAC Address of Intermediate System neighbour.
- b) Intermediate System Neighbors with variable length SNPA address

x CODE – 7
 x LENGTH – total length of the value field in octets
 x VALUE –

	No. of Octets
Lan Address Length	1
LAN Address (variable length)	LAN Address Length
LAN Address (variable length)	LAN Address Length

- **LAN Address Length** – length of SNPA address of Intermediate System neighbors.
- **LAN Address (variable length)** – variable length SNPA address of Intermediate System neighbors

This variable length field shall not be used for the Intermediate System Neighbor of which SNPA address length is 6 octets.

- **Padding** – This option may occur more than once. It is used to pad the PDU to at least *maxsize* - 1.
 - x CODE – 8.
 - x LENGTH – total length of the value field (may be zero).
 - x VALUE – LENGTH octets of arbitrary value.
- **Authentication Information** — information for performing authentication of the originator of the PDU.
 - x CODE — 10.
 - x LENGTH — variable from 1–254 octets
 - x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

- **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:
 - 0 — RESERVED
 - 1 — Cleartext Password
 - 2–254 — RESERVED
 - 255 — Routeing Domain private authentication method
- **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.6 Level 2 LAN IS to IS hello PDU

This PDU is multicast by Intermediate systems on broadcast circuits to the multi-destination address AllL2ISs. The purpose of this PDU is for Intermediate systems on broadcast circuits to discover the identity of other Level 2 Intermediate systems on that circuit. Trailing Pad options are inserted to make PDU Length equal to at least *maxsize* - 1 where

- dataLinkBlocksize
- originatingL2LSPBufferSize

(see 8.4.2).

	No. of Octets
Intradomain Routeing Protocol Discriminator	1
Length Indicator	1
Version/Protocol ID Extension	1
ID Length	1
R R R PDU Type	1
Version	1
Reserved	1
Maximum Area Addresses	1
Reserved/Circuit Type	1
Source ID	ID Length
Holding Time	2
PDU Length	2
R Priority	1
LAN ID	ID Length + 1
VARIABLE LENGTH FIELDS	VARIABLE

- Intradomain Routing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 16.

NOTE 56 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt.

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumAreaAddresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.

NOTE 57 This field is ignored on receipt.

- Reserved/Circuit Type – Most significant 6 bits reserved (Transmitted as zero, ignored on receipt). Low order bits (bits 1 and 2) indicate:
 - 0 – reserved value (if specified the entire PDU shall be ignored)
 - 1 – Level 1 only
 - 2 – Level 2 only (sender is Level 2 Intermediate System with **manualL2OnlyMode** set “True” for this circuit, and will use this link only for Level 2 traffic)
 - 3 – both Level 1 and Level 2 (sender is Level 2 Intermediate System, and will use this link both for Level 1 and Level 2 traffic)

NOTE 58 In a LAN Level 2 IIH PDU the Circuit Type shall be either 2 or 3

- Source ID – the system ID of transmitting Intermediate System
- Holding Time – Holding Timer to be used for this Intermediate System
- PDU Length – Entire length of this PDU, in octets, including header
- Reserved/Priority – Bit 8 reserved (Transmitted as zero, ignored on receipt). Bits 1 through 7 – priority for being LAN Level 2 Designated Intermediate System. Higher number has higher priority for being LAN Level 2 Designated Intermediate System. Unsigned integer.
- LAN ID – a field composed the system ID (1–8 octets) of the LAN Level 2 Designated Intermediate System, plus a low order octet assigned by LAN Level 2 Designated Intermediate System. Copied from LAN Level 2 Designated Intermediate System’s IIH PDU.

- VARIABLE LENGTH FIELDS – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received PDU that are not recognised shall be ignored.

Currently defined codes are:

- **Area Addresses** – the set of manualAreaAddresses of this Intermediate system.

x CODE – 1
 x LENGTH – total length of the value field.
 x VALUE –

	No. of Octets
Address Length	1
Area Address	Address Length
Address Length	1
Area Address	Address Length

· **Address Length** – Length of area address in octets.

· **Area Address** – Area address.

NOTE 59 This option field may be omitted, and is ignored on receipt if present.

- **Intermediate System Neighbours** – The set of Intermediate systems on this LAN to which adjacencies of neighborSystemType “L2 Intermediate System” exist in state “Up” or “Initialising” (i.e. those from which Level 2 IIH PDUs have been heard). This is permitted to appear more than once.

Two types of codes are defined:

- Intermediate System Neighbors with 6 octet MAC address

x CODE – 6
 x LENGTH – total length of the value field in octets
 x VALUE –

	No. of Octets
LAN Address	6
LAN Address	6

- **LAN Address (6 octet length)** – 6 octet MAC Address of Intermediate System neighbour.

- Intermediate System Neighbors with variable length SNPA address

x CODE – 7
 x LENGTH – total length of the value field in octets
 x VALUE –

	No. of Octets
Lan Address Length	1
LAN Address (variable length)	LAN Address Length
LAN Address (variable length)	LAN Address Length

- LAN Address Length – length of SNPA address of Intermediate System neighbors.
- LAN Address (variable length) – variable length SNPA address of Intermediate System neighbors

This variable length field shall not be used for the Intermediate System Neighbor of which SNPA address length is 6 octets.

- Padding – This option may occur more than once. It is used to pad the PDU to at least *maxsize* - 1.
 x CODE – 8.
 x LENGTH – total length of the value field (may be zero).
 x VALUE – LENGTH octets of arbitrary value.
- Authentication Information — information for performing authentication of the originator of the PDU.
 x CODE — 10.
 x LENGTH — variable from 1–254 octets
 x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

• Authentication Type — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED
 1 — Cleartext Password
 2–254 — RESERVED
 255 — Routing Domain private authentication method

• Authentication Value — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.7 Point-to-point IS to IS hello PDU

This PDU is transmitted by Intermediate systems on nonbroadcast circuits, after receiving an ISH PDU from the neighbour system. Its purpose is to determine whether the neighbour is a Level 1 or a Level 2 Intermediate System. Trailing pad options are inserted to make PDU Length equal to at least *maxsize* - 1 where *maxsize* is the maximum of

- dataLinkBlockSize
- originatingL1LSPBufferSize
- originatingL2LSPBufferSize

(see 8.2.3).

				No. of Octets
Intradomain Routing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
Reserved/Circuit Type				1
Source ID				ID Length
Holding Time				2
PDU Length				2
Local Circuit ID				1
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routing Protocol Discriminator — architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type — (bits 1 through 5) – 17.

NOTE 60 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt.

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter `maximumAreaAddresses`. This field shall take on one of the following values:
 - An integer between 1 and 254 inclusive, indicating a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for `maximumAreaAddresses`.
- Reserved/Circuit Type – Most significant 6 bits reserved (Transmitted as zero, ignored on receipt). Low order bits (bits 1 and 2) indicate:
 - 0 – reserved value (if specified the entire PDU shall be ignored)
 - 1 – Level 1 only
 - 2 – Level 2 only (sender is Level 2 Intermediate system with `manualL2OnlyMode` set “True” for this circuit, and will use this link only for Level 2 traffic)
 - 3 – both Level 1 and Level 2 (sender is Level 2 Intermediate system and will use this link both for Level 1 and Level 2 traffic)

- Source ID – the system ID of transmitting Intermediate system
- Holding Time – Holding Timer to be used for this Intermediate system
- PDU Length – Entire length of this PDU, in octets, including header
- Local Circuit ID – 1 octet unique ID assigned to this circuit when it is created by this Intermediate system.
- VARIABLE LENGTH FIELDS – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received PDU that are not recognised shall be ignored.

Currently defined codes are:

- Area Addresses – the set of manualAreaAddresses of this Intermediate system

x CODE – 1
x LENGTH – total length of the value field.
x VALUE –

	No. of Octets
Address Length	1
Area Address	Address Length
Address Length	1
Area Address	Address Length

· Address Length – Length of area address in octets.

· Area Address – Area address.

- Padding – This option may occur more than once. It is used to pad the PDU to at least *maxsize* - 1.
x CODE – 8.
x LENGTH – total length of the value field (may be zero).
x VALUE – LENGTH octets of arbitrary value.
- Authentication Information — information for performing authentication of the originator of the PDU.
x CODE – 10.
x LENGTH — variable from 1–254 octets
x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

· Authentication Type — a one octet identifier for the type of authentication to be carried out. The following values are defined:

- 0 — RESERVED
- 1 — Cleartext Password

2–254 — RESERVED

255 — Routeing Domain private authentication method

· **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.8 Level 1 link state PDU

Level 1 Link State PDUs are generated by Level 1 and Level 2 Intermediate systems, and propagated throughout an area. The contents of the Level 1 Link State PDU indicates the state of the adjacencies to neighbour Intermediate Systems, or pseudonodes, and End systems of the Intermediate system that originally generated the PDU.

				No. of Octets
Intradomain Routeing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Remaining Lifetime				2
LSP ID				ID Length + 2
Sequence Number				4
Checksum				2
P	ATT	LSPDBOL	IS Type	1
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routeing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension — 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routeing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 18.
- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumArea-Addresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.

- The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.

- **PDU Length** – Entire Length of this PDU, in octets, including header.
- **Remaining Lifetime** – Number of seconds before LSP considered expired
- **LSP ID** – the system ID of the source of the Link State PDU. It is structured as follows:

	No. of Octets
Source ID	ID Length
Pseudonode ID	1
LSP Number	1

- **Sequence Number** – sequence number of LSP
- **Checksum** – Checksum of contents of LSP from **Source ID** to end. Checksum is computed as described in 7.3.11.
- **P/ATT/LSPDBOL/IS Type**
 - **P** – Bit 8, indicates when set that the issuing Intermediate System supports the Partition Repair optional function.
 - **ATT** - Bits 7-4 indicate, when set, that the issuing Intermediate System is 'attached' to other areas using:

x Bit 4 - the Default Metric

x Bit 5 - the Delay Metric

x Bit 6 - the Expense Metric

x Bit 7 - the Error Metric.

- **LSPDBOL** – Bit 3 – A value of 0 indicates no LSP Database Overload, and a value of 1 indicates that the LSP Database is Overloaded. An LSP with this bit set will not be used by any decision process to calculate routes to another IS through the originating system.
- **IS Type** – Bits 1 and 2 indicate the type of Intermediate System – One of the following values:

x 0 – Unused value

x 1 – (i.e. bit 1 set) Level 1 Intermediate system

x 2 – Unused value

x 3 – (i.e. bits 1 and 2 set) Level 2 Intermediate system.

- **VARIABLE LENGTH FIELDS** – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received LSP that are not recognised are ignored and passed through unchanged.

Currently defined codes are:

- **Area Addresses** – the set of **manualAreaAddresses** of this Intermediate system. For non-pseudonode LSPs this option shall always be present in the LSP with LSP number zero, and shall never be present in an LSP with non-zero LSP

number. It shall appear before any Intermediate System Neighbours or End System Neighbours options. This option shall never be present in pseudonode LSPs.

- x CODE – 1
- x LENGTH – total length of the value field.
- x VALUE –

	No. of Octets
Address Length	1
Area Address	Address Length
Address Length	1
Area Address	Address Length

· Address Length – Length of area address in octets.

· Area Address – Area address.

- originatingLSPBufferSize – the local value for originatingLSPBufferSize

This may appear only once in an LSP with any LSP number. It shall appear before any Intermediate System Neighbours or End System Neighbours options.

NOTE 61 It is recommended that this option be used only in the LSP with LSP number zero. However, the presence or absence of this option in any LSP shall not be required.

- x CODE – 14
- x LENGTH – 2
- x VALUE – 512-1492

- Intermediate System Neighbours – Intermediate system and pseudonode neighbours.

This is permitted to appear more than once, and in an LSP with any LSP number. However, all the Intermediate System Neighbours options shall precede the End System Neighbours options. i.e. they shall appear before any End System Neighbours options in the same LSP and no End System Neighbours options shall appear in an LSP with lower LSP number.

- x CODE – 2.
- x LENGTH – 1 plus a multiple of (IDLength+5).
- x VALUE –

	No. of Octets
Virtual Flag	1
0 I/E Default Metric	1
S I/E Delay Metric	1
S I/E Expense Metric	1
S I/E Error Metric	1
Neighbor ID	ID Length + 1
0 I/E Default Metric	1
S I/E Delay Metric	1
S I/E Expense Metric	1
S I/E Error Metric	1
Neighbor ID	ID Length +1

· Virtual Flag is a Boolean. If equal to 1, this indicates the link is really a Level 2 path to repair an area partition. (Level 1 Intermediate Systems would always report this octet as 0 to all neighbours).

· **Default Metric** is the value of the default metric for the link to the listed neighbour. Bit 8 of this field is reserved. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Delay Metric** is the value of the delay metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Expense Metric** is the value of the expense metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Error Metric** is the value of the error metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Neighbour ID.** For Intermediate System neighbours, neighbour ID field consists of the neighbour system’s ID, followed by an octet containing the value zero. For pseudonode neighbours, the first ID Length octets is the LAN Level 1 Designated Intermediate System’s ID, and the last octet is a non-zero quantity defined by the LAN Level 1 Designated Intermediate System.

• **End System Neighbours** – End system neighbours This may appear more than once, and in an LSP with any LSP number. See the description of the **Intermediate System Neighbours** option above for the relative ordering constraints. Only adjacencies with identical costs can appear in the same list.

- x CODE – 3.
- x LENGTH – 4, plus a multiple of IDLength.
- x VALUE –

Virtual Flag			No. of Octets
0	I/E	Default Metric	1
S	I/E	Delay Metric	1
S	I/E	Expense Metric	1
S	I/E	Error Metric	1
Neighbor ID			ID Length
Neighbor ID			ID Length

· **Default Metric** is the value of the default metric for the link to each of the listed neighbours. Bit 8 of this field is reserved. Bit 7 (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an internal metric.

· **Delay Metric** is the value of the delay metric for the link to each of the listed neighbours. If this IS does not support this metric it shall set the bit “S” to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an internal metric.

· **Expense Metric** is the value of the expense metric for the link to each of the listed neighbours. If this IS does not support this metric it shall set the bit “S” to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an internal metric.

· **Error Metric** is the value of the error metric for the link to each of the listed neighbour. If this IS does not support this metric it shall set the bit “S” to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an internal metric.

· **Neighbour ID** – system ID of End system neighbour.

• **Authentication Information** — information for performing authentication of the originator of the PDU.

- x CODE — 10.
- x LENGTH — variable from 1–254 octets
- x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

· **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED

1 — Cleartext Password

2–254 — RESERVED

255 — Routeing Domain private authentication method

· **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.9 Level 2 link state PDU

Level 2 Link State PDUs are generated by Level 2 Intermediate systems, and propagated throughout the level 2 domain. The contents of the Level 2 Link State PDU indicates the state of the adjacencies to neighbour Level 2 Intermediate Systems, or pseudonodes, and to reachable address prefixes of the Intermediate system that originally generated the PDU.

	No. of Octets
Intradomain Routeing Protocol Discriminator	1
Length Indicator	1
Version/Protocol ID Extension	1
ID Length	1
R R R PDU Type	1
Version	1
Reserved	1
Maximum Area Addresses	1
PDU Length	2
Remaining Lifetime	2
LSP ID	ID Length + 2
Sequence Number	4
Checksum	2
P ATT LSPDBOL IS Type	1
VARIABLE LENGTH FIELDS	VARIABLE

- Intradomain Routeing Protocol Discriminator – architectural constant
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routeing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 20.

NOTE 62 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt.

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumAreaAddresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.

NOTE 63 This field is ignored on receipt.

- PDU Length – Entire Length of this PDU, in octets, including header.
- Remaining Lifetime – Number of seconds before LSP considered expired
- LSP ID – the system ID of the source of the Link State PDU. It is structured as follows:

	No. of Octets
Source ID	ID Length
Pseudonode ID	1
LSP Number	1

- Sequence Number – sequence number of LSP
- Checksum – Checksum of contents of LSP from Source ID to end. Checksum is computed as described in 7.3.11.
- P/ATT/LSPDBOL/IS Type
 - P – Bit 8, indicates when set that the issuing Intermediate System supports the Partition Repair optional function.
 - ATT - Bits 7-4 indicate, when set, that the issuing Intermediate System is ‘attached’ to other areas using:

NOTE 64 These bits are ignored on receipt.

- x Bit 4 - the Default Metric
- x Bit 5 - the Delay Metric
- x Bit 6 - the Expense Metric
- x Bit 7 - the Error Metric.

- LSPDBOL – Bit 3 – A value of 0 indicates no LSP Database Overload, and a value of 1 indicates that the LSP Database is Overloaded. An LSP with this bit set will not be used by any decision process to calculate routes to another IS through the originating system.

- IS Type – Bits 1 and 2 indicate the type of Intermediate System – One of the following values:

- x 0 – Unused value
- x 1 – (i.e. bit 1 set) Level 1 Intermediate system
- x 2 – Unused value
- x 3 – (i.e. bits 1 and 2 set) Level 2 Intermediate system.

NOTE 65 In a Level 2 Link State PDU, IS Type shall be 3.

- VARIABLE LENGTH FIELDS – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received LSP that are not recognised are ignored and passed through unchanged.

Currently defined codes are:

- **Area Addresses** – the set of **partitionArea-Addresses** of this Intermediate system, if the system supports partition repair, otherwise the set of **areaAddresses** of the IS. For non-pseudonode LSPs this option shall always be present in the LSP with LSP number zero, and shall never be present in an LSP with non-zero LSP number. It shall appear before any **Intermediate System Neighbours** or **Prefix Neighbours** options. This option shall never be present in pseudonode LSPs.

x CODE – 1
 x LENGTH – total length of the value field.
 x VALUE –

	No. of Octets
Address Length	1
Area Address	Address Length
Address Length	1
Area Address	Address Length

· **Address Length** – Length of area address in octets.

· **Area Address** – Area address.

- **Partition Designated Level 2 Intermediate System** – ID of Designated Level 2 Intermediate System for the partition. For non-pseudonode LSPs issued by Intermediate Systems which support the partition repair optional function and which are currently ATTACHED, this option shall always be present in the LSP with LSP number zero, and shall never be present in an LSP with non-zero LSP number. It shall appear before any **Intermediate System Neighbours** or **Prefix Neighbours** options. This option shall never be present in pseudonode LSPs.

x CODE – 4.
 x LENGTH – IDLength
 x VALUE – systemID of Partition Designated Level 2 Intermediate System for the partition.

- **originatingLSPBufferSize** – the local value for originatingL2LSPBufferSize

This may appear only once in an LSP with any LSP number. It shall appear before any **Intermediate System Neighbours** or **Prefix Neighbours** options.

NOTE 66 It is recommended that this option be used only in the LSP with LSP number zero. However, the presence or absence of this option in any LSP shall not be required.

x CODE – 14
 x LENGTH – 2
 x VALUE – 512 – 1492

- **Intermediate System Neighbours** – Intermediate system and pseudonode neighbours.

This is permitted to appear more than once, and in an LSP with any LSP number. However, all the **Intermediate System Neighbours** options shall precede the **Prefix Neighbours** options. i.e. they shall appear before any **Prefix Neighbour** options in the same LSP and no **Prefix Neighbour** options shall appear in an LSP with lower LSP number.

- x CODE – 2.
- x LENGTH – 1 plus a multiple of (IDLength+5).
- x VALUE –

Virtual Flag			No. of Octets
0	I/E	Default Metric	1
S	I/E	Delay Metric	1
S	I/E	Expense Metric	1
S	I/E	Error Metric	1
Neighbor ID			ID Length + 1
0	I/E	Default Metric	1
S	I/E	Delay Metric	1
S	I/E	Expense Metric	1
S	I/E	Error Metric	1
Neighbor ID			ID Length + 1

· **Virtual Flag** is a Boolean. If equal to 1, this indicates the link is really a Level 2 path to repair an area partition. (Level 1 Intermediate Systems would always report this octet as 0 to all neighbours).

· **Default Metric** is the value of the default metric for the link to the listed neighbour. Bit 8 of this field is reserved. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Delay Metric** is the value of the delay metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Expense Metric** is the value of the expense metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Error Metric** is the value of the error metric for the link to the listed neighbour. If this IS does not support this metric it shall set bit “S” to 1 to indicate that the metric is unsupported. Bit 7 of this field (marked I/E) indicates the metric type, and shall contain the value “0”, indicating an Internal metric.

· **Neighbour ID** – For Intermediate System neighbours, the first ID Length octets are the neighbour’s system ID, and the last octet is 0. For pseudonode neighbours, the first ID Length octets is the LAN Level 2 Designated Intermediate System’s ID, and the last octet is a non-zero quantity defined by the LAN Level 2 Designated Intermediate System.

- **Prefix Neighbours** – reachable address prefix neighbours

This may appear more than once, and in an LSP with any LSP number. See the description of the **Intermediate System Neighbours** option above for the relative ordering constraints. Only adjacencies with identical costs can appear in the same list.

- x CODE – 5.
- x LENGTH – Total length of the VALUE field.
- x VALUE –

			No. of Octets
0	I/E	Default Metric	1
S	I/E	Delay Metric	1
S	I/E	Expense Metric	1
S	I/E	Error Metric	1
Address Prefix Length			1
Address Prefix			$\lceil \text{Address Prefix Length} / 2 \rceil$
Address Prefix Length			1
Address Prefix			$\lceil \text{Address Prefix Length} / 2 \rceil$

· **Default Metric** is the value of the default metric for the link to each of the listed neighbours. Bit 8 of this field is reserved. Bit 7 (marked I/E) indicates the metric type, and may be set to zero indicating an internal metric, or may be set to 1 indicating an external metric.

· **Delay Metric** is the value of the delay metric for the link to each of the listed neighbours. If this IS does not support this metric it shall set the bit "S" to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and may be set to zero indicating an internal metric, or may be set to 1 indicating an external metric.

· **Expense Metric** is the value of the expense metric for the link to each of the listed neighbours. If this IS does not support this metric it shall set the bit "S" to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and may be set to zero indicating an internal metric, or may be set to 1 indicating an external metric.

· **Error Metric** is the value of the error metric for the link to each of the listed neighbour. If this IS does not support this metric it shall set the bit "S" to 1 to indicate that the metric is unsupported. Bit 7 (marked I/E) indicates the metric type, and may be set to zero indicating an internal metric, or may be set to 1 indicating an external metric.

· **Address Prefix Length** is the length in semi-octets of the following prefix. A length of zero indicates a prefix that matches all NSAPs.

· **Address Prefix** is a reachable address prefix encoded as described in 7.1.6. If the length in semi-octets is odd, the prefix is padded out to an integral number of octets with a trailing zero semi-octet.

Note that the area addresses listed in the **Area Addresses** option field of Level 2 Link State PDU with LSP number zero, are understood to be reachable address neighbours with cost zero. They are not listed separately in the **Prefix Neighbours** options.

- **Authentication Information** — information for performing authentication of the originator of the PDU.

x CODE — 10

x LENGTH — variable from 1–254 octets

x VALUE —

		No. of Octets
Authentication Type		1
Authentication Value		VARIABLE

· **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED

1 — Cleartext Password

2–254 — RESERVED

255 — Routeing Domain private authentication method

· **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.10 Level 1 complete sequence numbers PDU

				No. of Octets
Intradomain Routeing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Source ID				ID Length + 1
Start LSP ID				ID Length + 2
End LSP ID				ID Length + 2
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routeing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routeing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 24.

NOTE 67 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt.

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumArea-Addresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.
- PDU Length – Entire Length of this PDU, in octets, including header
- Source ID – the system ID of Intermediate System (with zero Circuit ID) generating this Sequence Numbers PDU.
- Start LSP ID – the LSP ID of first LSP in the range covered by this Complete Sequence Numbers PDU.
- End LSP ID – the LSP ID of last LSP in the range covered by this Complete Sequence Numbers PDU.
- VARIABLE LENGTH FIELDS – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received CSNP that are not recognised are ignored.

Currently defined codes are:

- **LSP Entries** – This may appear more than once. The option fields, if they appear more than once, shall appear sorted into ascending LSPID order.

x CODE – 9

x LENGTH – total length of the value field.

x VALUE – a list of LSP entries of the form:

	No. of Octets
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2

- **Remaining Lifetime** – Remaining Lifetime of LSP.
- **LSP ID** – system ID of the LSP to which this entry refers.
- **LSP Sequence Number** – Sequence number of LSP.
- **Checksum** – Checksum reported in LSP.

The entries shall be sorted into ascending LSPID order (the LSP number octet of the LSPID is the least significant octet).

- **Authentication Information** — information for performing authentication of the originator of the PDU.

x CODE — 10.

x LENGTH — variable from 1–254 octets

x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

- **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED

1 — Cleartext Password

2–254 — RESERVED

255 — Routeing Domain private authentication method

- **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.11 Level 2 complete sequence numbers PDU

				No. of Octets
Intradomain Routing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Source ID				ID Length + 1
Start LSP ID				ID Length + 2
End LSP ID				ID Length + 2
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 25.

NOTE 68 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter **maximumArea-Addresses**. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for **maximumAreaAddresses**.

NOTE 69 This field is ignored on receipt.

- PDU Length – Entire Length of this PDU, in octets, including header
- Source ID – the system ID of Intermediate System (with zero Circuit ID) generating this Sequence Numbers PDU.
- Start LSP ID – the LSP ID of first LSP in the range covered by this Complete Sequence Numbers PDU.
- End LSP ID – the LSP ID of last LSP in the range covered by this Complete Sequence Numbers PDU.

- VARIABLE LENGTH FIELDS – fields of the form:

	No. of Octets
CODE	1
LENGTH	1
VALUE	LENGTH

Any codes in a received CSNP that are not recognised are ignored.

Currently defined codes are:

- **LSP Entries** – this may appear more than once. The option fields, if they appear more than once, shall appear sorted into ascending LSPID order.

- x CODE – 9
- x LENGTH – total length of the value field.
- x VALUE – a list of LSP entries of the form:

	No. of Octets
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2

- Remaining Lifetime – Remaining Lifetime of LSP.
- LSP ID – the system ID of the LSP to which this entry refers.
- LSP Sequence Number – Sequence number of LSP.
- Checksum – Checksum reported in LSP.

The entries shall be sorted into ascending LSPID order (the LSP number octet of the LSPID is the least significant octet).

- **Authentication Information** — information for performing authentication of the originator of the PDU.

- x CODE — 10.
- x LENGTH — variable from 1–254 octets
- x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

- **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

- 0 — RESERVED
- 1 — Cleartext Password
- 2–254 — RESERVED
- 255 — Routing Domain private authentication method

- **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.12 Level 1 partial sequence numbers PDU

				No. of Octets
Intradomain Routeing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Source ID				ID Length + 1
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routeing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routeing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 26.

NOTE 70 Bits 6, 7 and 8 are Reserved, which means they are transmitted as 0 and ignored on receipt.

- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter `maximumAreaAddresses`. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for `maximumAreaAddresses`.
- PDU Length – Entire Length of this PDU, in octets, including header
- Source ID – the system ID of Intermediate system (with zero Circuit ID) generating this Sequence Numbers PDU.
- VARIABLE LENGTH FIELDS – fields of the form:

		No. of Octets
CODE		1
LENGTH		1
VALUE		LENGTH

Any codes in a received PSNP that are not recognised are ignored.

Currently defined codes are:

- **LSP Entries** – this may appear more than once. The option fields, if they appear more than once, shall appear sorted into ascending LSPID order.

x CODE – 9
 x LENGTH – total length of the value field.
 x VALUE – a list of LSP entries of the form:

	No. of Octets
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2

- **Remaining Lifetime** – Remaining Lifetime of LSP.
- **LSP ID** – the system ID of the LSP to which this entry refers.
- **LSP Sequence Number** – Sequence number of LSP.
- **Checksum** – Checksum reported in LSP.

The entries shall be sorted into ascending LSPID order (the LSP number octet of the LSPID is the least significant octet).

- **Authentication Information** — information for performing authentication of the originator of the PDU.

x CODE — 10.
 x LENGTH — variable from 1–254 octets
 x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

- **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED
 1 — Cleartext Password
 2–254 — RESERVED
 255 — Routing Domain private authentication method

- **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

9.13 Level 2 partial sequence numbers PDU

				No. of Octets
Intradomain Routeing Protocol Discriminator				1
Length Indicator				1
Version/Protocol ID Extension				1
ID Length				1
R	R	R	PDU Type	1
Version				1
Reserved				1
Maximum Area Addresses				1
PDU Length				2
Source ID				ID Length + 1
VARIABLE LENGTH FIELDS				VARIABLE

- Intradomain Routeing Protocol Discriminator – architectural constant (see table 2)
- Length Indicator – Length of fixed header in octets
- Version/Protocol ID Extension – 1
- ID Length — Length of the ID field of NSAP addresses and NETs used in this routeing domain. This field shall take on one of the following values:
 - An integer between 1 and 8, inclusive, indicating an ID field of the corresponding length
 - The value zero, which indicates a 6 octet ID field length
 - The value 255, which means a null ID field (i.e. zero length)

All other values are illegal and shall not be used.

- PDU Type (bits 1 through 5) – 27.
- Version – 1
- Reserved — transmitted as zero, ignored on receipt
- Maximum Area Addresses — number of area addresses permitted for this ISs area, as derived from the value of the System Management parameter `maximumAreaAddresses`. This field shall take on of the following values:
 - An integer between 1 and 254 inclusive, indicated a corresponding number of area addresses supported.
 - The value zero, which is treated upon reception as if it were equal to three, and which the IS may use if it supports only a value of 3 for `maximumAreaAddresses`.

NOTE 71 This field is ignored on receipt.

- PDU Length – Entire Length of this PDU, in octets, including header
- Source ID – the system ID of Intermediate system (with zero Circuit ID) generating this Sequence Numbers PDU.
- VARIABLE LENGTH FIELDS – fields of the form:

		No. of Octets
CODE		1
LENGTH		1
VALUE		LENGTH

Any codes in a received PSNP that are not recognised are ignored.

Currently defined codes are:

- **LSP Entries** – this may appear more than once. The option fields, if they appear more than once, shall appear sorted into ascending LSPID order.

x CODE – 9

x LENGTH – total length of the value field.

x VALUE – a list of LSP entries of the form:

	No. of Octets
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2
Remaining Lifetime	2
LSP ID	ID Length + 2
LSP Sequence Number	4
Checksum	2

· **Remaining Lifetime** – Remaining Lifetime of LSP.

· **LSP ID** – the system ID of the LSP to which this entry refers.

· **LSP Sequence Number** – Sequence number of LSP.

· **Checksum** – Checksum reported in LSP. The entries shall be sorted into ascending LSPID order (the LSP number octet of the LSPID is the least significant octet).

- **Authentication Information** — information for performing authentication of the originator of the PDU.

x CODE — 10.

x LENGTH — variable from 1–254 octets

x VALUE —

	No. of Octets
Authentication Type	1
Authentication Value	VARIABLE

· **Authentication Type** — a one octet identifier for the type of authentication to be carried out. The following values are defined:

0 — RESERVED

1 — Cleartext Password

2–254 — RESERVED

255 — Routing Domain private authentication method

· **Authentication Value** — determined by the value of the authentication type. If Cleartext Password as defined in this International Standard is used, then the authentication value is an octet string.

10 System environment

10.1 Generating jitter on timers

When PDUs are transmitted as a result of timer expiration, there is a danger that the timers of individual systems may become synchronised. The result of this is that the traffic distribution will contain peaks. Where there are a large number of synchronised systems, this can cause overloading of both the transmission medium and the systems receiving the PDUs. In order to prevent this from occurring, all periodic timers, the expiration of which can cause the transmission of PDUs, shall have “jitter” introduced as defined in the following algorithm.

CONSTANT

Jitter = 25;

(* The percentage jitter as defined in the architectural constant Jitter *) (see table 2)

Resolution = 100;

(* The timer resolution in ms *)

PROCEDURE Random(max : Integer): Integer;

(* This procedure delivers a Uniformly distributed random integer R such that $0 < R < \text{max}$ *)

PROCEDURE WaitUntil(time: Integer)

(* This procedure waits the specified number of ms and then returns *)

PROCEDURE CurrentTime(): Integer

(* This procedure returns the current time in ms*)

PROCEDURE

DefineJitteredTimer(baseTimeValueInSeconds:

Integer; expirationAction : Procedure);

VAR

baseTimeValue, maximumTimeModifier, waitTime

: Integer;

nextexpiration : Time;

BEGIN

baseTimeValue := baseTimeValueInSeconds * 1000

/ Resolution;

maximumTimeModifier := baseTimeValue * Jitter /

100; (* Compute maximum possible jitter *)

WHILE running **DO**

BEGIN

(* First compute next expiration time *)

randomTimeModifier :=

Random(maximumTimeModifier);

waitTime := baseTimeValue - randomTimeModifier;

nextexpiration := CurrentTime() + waitTime;

(* Then perform expiration Action *)

expirationAction;

WaitUntil(nextexpiration);

END (* of Loop *)

END (* of DefineJitteredTimer *)

Thus the call "DefineJitteredTimer(HelloTime, SendHello-PDU);", where "HelloTime" is 10 s, will cause the action "SendHelloPDU" to be performed at random intervals of between 7,5 and 10 s. The essential point of this algorithm is that the value of "randomTimeModifier" is randomised within the inner loop. Note that the new expiration time is set immediately on expiration of the last interval, rather than when the expiration action has been completed.

The time resolution shall be less than or equal to 100 ms. It is recommended to be less than or equal to 10 ms. The time resolution is the maximum interval that can elapse without there being any change in the value of the timer. The periodic transmission period shall be random or pseudo-random in the specified range, with uniform distribution across similar implementations.

10.2 Resolution of timers

All timers specified in units of seconds shall have a resolution of no less than ± 1 s.

All timers specified in units of ms shall have a resolution of no less than ± 10 ms.

10.3 Requirements on the operation of ISO 9542

This International Standard places certain requirements on the use of ISO 9542 by Intermediate systems which go beyond those mandatory requirements stated in the conformance clause of ISO 9542. These requirements are

- a) The IS shall operate the Configuration Information functions on all types of subnetworks supported by the IS. This includes:
 - 1) the reception of ESH PDUs on all types of circuits;
 - 2) the transmission of ISH PDUs on broadcast circuits; and
 - 3) the transmission and reception of ISH PDUs on point-to-point circuits.
- b) The IS shall enable the “All Intermediate Systems” multi-destination subnetwork address.
- c) When sending ISH PDUs, or redirecting an End system to a neighbour Intermediate system, as described in 7.4.3.3, the numerically lowest Network Entity title shall be chosen as the NET placed in the ISO 9542 PDU. This minimises memory usage in the End systems by ensuring that all ISs identify themselves and each other to ESs using the same Network Entity titles.

10.4 Requirements on the operation of ISO 8473

If the area partition repair functions described in 7.2.10 are implemented by the Intermediate system, then the IS requires a number of End system functions of ISO 8473 (such as NPDU reassembly) as well as the Intermediate system functions. Such ISs shall therefore conform to both the End system and Intermediate system requirements of ISO 8473

11 System management

11.1 General

The operation of the Intra-domain IS-IS routing functions may be monitored and controlled using System Management. This clause is the management specification for this International Standard in the GDMO notation as defined in ISO/IEC 10165-4. Generic managed object and attribute definitions are imported from ISO/IEC 10165-5 and from ISO/IEC 10733.

The containment hierarchy for this International Standard is illustrated below in figure 8.

Annex E contains a number of definitions which are intended to be incorporated in ISO/IEC 10165-5. In those cases where a GDMO imported definition cannot be found by consulting ISO/IEC 10165-5, or conflicts with the definition in Annex E, the definition in Annex E shall be used.




Figure 8 - Containment and naming hierarchy

--11.2 GDMO definition

--11.2.1 Common GDMO definitions

--11.2.1.1 Behaviours

supplyValueOnCreate-B BEHAVIOUR

DEFINED AS

Value is supplied either by the protocol machine if the MO is automatically created, or supplied via the CREATE Operation. The value cannot be changed thereafter.;

resettingTimer-B BEHAVIOUR

DEFINED AS

This attribute specifies the interval between certain events in the operation of the protocol state machine. If the value of this attribute is changed to a new value while the protocol state machine is in operation, the implementation shall take the necessary steps to ensure that for any time interval which was in progress when the corresponding attribute was changed, the next expiration of the that interval takes place the specified time after the original start of that interval, or immediately, whichever is later. The precision with which this time shall be implemented shall be the same as that associated with the basic operation of the timer attribute;

resettingHoldingTimer-B BEHAVIOUR

DEFINED AS

This attribute specifies the interval between certain events in the operation of the protocol state machine. If the value of this attribute is changed to a new value while the protocol state machine is in operation, the implementation shall take the necessary steps to ensure that for any time interval which was in progress when the corresponding attribute was changed, the next expiration of that interval takes place no later than the expiration of the interval in progress or the specified interval, whichever is sooner. The precision with which this time shall be implemented shall be the same as that associated with the basic operation of the timer attribute;

replaceOnlyWhileDisabled-B BEHAVIOUR**DEFINED AS**

This attribute shall only permit the REPLACE operation to be performed on it while the MO is in the Disabled Operational State. An attempt to perform a REPLACE operation while the MO is in the Enabled Operation State shall fail with the generation of the constraintViolation specific error.;

resourceLimiting-B BEHAVIOUR**DEFINED AS**

This attribute places limits on some "resource". In general implementations may allocate resources up to this limit when the managed object is enabled and it may be impossible to change the allocation without first disabling and re-enabling the managed object. Therefore this International Standard only requires that it shall be possible to perform a REPLACE operation on this attribute while the MO is disabled. However some implementations may be able to change the allocation of resources without first disabling the MO. In this case it is permitted to increase the value of the attribute at any time, but it shall not be decreased below the currently "used" value of the resource. Where an attempt to perform a REPLACE operation fails either because the MO is enabled, or because an attempt has been made to decrease the value, the REPLACE operation shall fail with the generation of the constraintViolation specific error.;

--11.2.1.2 Common attributes**addressPrefix ATTRIBUTE**

WITH ATTRIBUTE SYNTAX ISIS.AddressPrefix;

MATCHES FOR EQUALITY, SUBSTRINGS;

BEHAVIOUR addressPrefix-B BEHAVIOUR

DEFINED AS An Area Address of a destination area or a prefix for a reachable address;;

REGISTERED AS {ISIS.aoi addressPrefix (98)};

authenticationFailures ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR authenticationFailures-B BEHAVIOUR

DEFINED AS Count of authentication Failure events;;

REGISTERED AS {ISIS.aoi authenticationFailures (117)};

networkEntityTitle ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.NAddress;

MATCHES FOR EQUALITY;

REGISTERED AS {ISIS.aoi networkEntityTitle (88)};

--11.2.1.3 Common parameters**constraintViolation PARAMETER**

CONTEXT SPECIFIC-ERROR;

WITH SYNTAX ISIS.Null;

BEHAVIOUR constraintViolation-B BEHAVIOUR

DEFINED AS

The specific error returned on failure of a REPLACE operation when the MO prohibits such operations under certain conditions, for example while the MO is in the disabled operational state;;

REGISTERED AS {ISIS.proi constraintViolation (10)};

notificationReceivingAdjacency PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.LocalDistinguishedName;

BEHAVIOUR notificationReceivingAdjacency-B BEHAVIOUR

DEFINED AS The local managed object name of the adjacency upon which the NPDU was received;;

REGISTERED AS {ISIS.proi notificationReceivingAdjacency (11)};

notificationIDLength PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.IDLength;

BEHAVIOUR notificationIDLength-B BEHAVIOUR

DEFINED AS The IDLength specified in the ignored PDU;;

REGISTERED AS {ISIS.proi notificationIDLength (12)};

notificationAreaAddress PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.AreaAddress;

BEHAVIOUR notificationAreaAddress-B BEHAVIOUR

DEFINED AS The Area Address which caused MaximumAreaAddresses to be exceeded;;

REGISTERED AS {ISIS.proi notificationAreaAddress (13)};

notificationAreaAddresses PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.AreaAddresses;

BEHAVIOUR notificationAreaAddresses-B BEHAVIOUR

DEFINED AS The set of Area Address from the neighbour's IIH PDU;;

REGISTERED AS {ISIS.proi notificationAreaAddresses (29)};

notificationSourceId PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.SourceId;

BEHAVIOUR notificationSourceId-B BEHAVIOUR

DEFINED AS The source ID of the LSP;;

REGISTERED AS {ISIS.proi notificationSourceId (14)};

notificationMaximumAreaAddresses PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.MaximumAreaAddresses;

BEHAVIOUR notificationMaximumAreaAddresses-B BEHAVIOUR

DEFINED AS The maximum area addresses field of the received PDU;;

REGISTERED AS {ISIS.proi notificationMaximumAreaAddresses (28)};

notificationVirtualLinkChange PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.VirtualLinkChange;

BEHAVIOUR notificationVirtualLinkChange-B BEHAVIOUR

DEFINED AS

This indicates whether the event was generated as a result of the creation or deletion of a Virtual Link between two Level 2 Intermediate Systems;;

REGISTERED AS {ISIS.proi notificationVirtualLinkChange (15)};

notificationVirtualLinkAddress PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.NAddress;

BEHAVIOUR notificationVirtualLinkAddress-B BEHAVIOUR

DEFINED AS

The Network Entity Title of the Level 2 Intermediate System at the remote end of the virtual link;;

REGISTERED AS {ISIS.proi notificationVirtualLinkAddress (16)};

notificationSystemId PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.SystemId;

BEHAVIOUR notificationSystemId-B BEHAVIOUR

DEFINED AS The system ID of the source system generating the notification;;

REGISTERED AS {ISIS.proi notificationSystemId (19)};

notificationVersion PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.Version;

BEHAVIOUR notificationVersion-B BEHAVIOUR

DEFINED AS The version number reported by the other system;;
 REGISTERED AS {ISIS.proi notificationVersion (23)};;

notificationDesignatedIntermediateSystemChange PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.DesignatedISChange;

BEHAVIOUR notificationDesignatedIntermediateSystemChange-B BEHAVIOUR

DEFINED AS The direction of the change in Designated Intermediate System status of this system;;

REGISTERED AS {ISIS.proi notificationDesignatedIntermediateSystemChange (24)};;

notificationOverloadStateChange PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.OverloadStateChange;

BEHAVIOUR notificationOverloadStateChange-B BEHAVIOUR

DEFINED AS The direction of the change in Overload status;;

REGISTERED AS {ISIS.proi notificationOverloadStateChange (25)};;

reservedName PARAMETER

CONTEXT SPECIFIC-ERROR;

WITH SYNTAX ISIS.Null;

BEHAVIOUR reservedName-B BEHAVIOUR

DEFINED AS

The specific error returned on failure of a create operation using a name reserved for other purposes.;;

REGISTERED AS {ISIS.proi reservedName (27)};;

notificationLSPHeader PARAMETER

CONTEXT EVENT-INFO;

WITH SYNTAX ISIS.OctetString;

BEHAVIOUR notificationLSPHeader-B BEHAVIOUR

DEFINED AS The header of the LSP;;

REGISTERED AS {ISIS.proi notificationLSPHeader(30)};;

--11.2.2 Conditional packages attached to objects defined in ISO/IEC 10733

--11.2.2.1 Basis ISIS package attached to cLNS managed object

cLNSISISBasic-P PACKAGE

BEHAVIOUR cLNSISISBasicImportedAlarmNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsAlarm notification from ISO/IEC 10165-2. It is used to report the following protocol events:

pDUDiscard:

Generated when an 8473 data PDU is discarded. In addition to the parameters specified by ISO/IEC 10733, the local adjacencyId of the adjacency upon which the NPDU was received shall be reported in the additionalInformation field using the notificationReceivingAdjacency parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it.

ISPL1DatabaseOverload:

Generated when the l1State of the system changes between On and Waiting or Waiting and On. The resulting state, shall be reported in the additionalInformation field using the notificationOverload-StateChange parameter and in the case of Waiting the SourceId of the LSP which precipitated the overload shall be reported in the additionalInformation field using the notificationSourceId parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.ISPL1DatabaseOverload shall be reported in the specificProblems parameter. The probable-Cause shall be set to NLM.resourceAtOrNearingCapacity. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters

in the additionalInformation field. The occurrence of this event shall be counted by the ISPL1Database-Overloads counter.

manualAddressDroppedFromArea:

Generated when one of the manualAreaAddresses assigned to this system is ignored when computing partitionAreaAddresses or areaAddresses because there are more than MaximumAreaAddresses distinct Area Addresses. The ignored Area Address shall be reported in the additionalInformation field using the notificationAreaAddress parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The event is generated once for each Area Address in manualAreaAddresses which is dropped. It is not generated again for that Area Address until after it has been reinstated into areaAddresses (i.e. it is only the action of dropping the Area Address and not the state of being dropped, which causes the event to be generated). The value ISIS.manualAddressDroppedFromArea shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additional-Information field. The occurrence of this event shall be counted by the manualAddressDroppedFromAreas counter.

corruptedLSPDetected:

Generated when a corrupted Link State PDU is detected in memory. The value ISIS.corruptedLSPDetected shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.corruptData. The perceivedSeverity shall be set to 'Minor'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the corruptedLSPsDetected counter.

attemptToExceedMaximumSequenceNumber:

Generated when an attempt is made to increment the sequence number of an LSP beyond the maximum sequence number. The value ISIS.attemptToExceedMaximumSequenceNumber shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.communicationsProtocolError. The perceivedSeverity shall be set to 'Major'. A subsequent communications-Alarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the attemptsToExceedMaximumSequenceNumber counter.

iDFieldLengthMismatch:

Generated when a PDU is received with a different value for ID field length to that of the receiving Intermediate system. The received ID LENGTH and SOURCE ID are reported in the additionalInformation field using the notificationIDLength and notificationSourceId parameters respectively. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.iDFieldLengthMismatch shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the iDFieldLengthMismatches counter.

maximumAreaAddressesMismatch:

Generated when a PDU is received with a different value for maximumAreaAddresses from that of the receiving Intermediate system. The received Maximum Area Addresses and Source ID are reported in the additionalInformation field using the notificationMaximumAreaAddresses and notificationSourceId parameters respectively. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.MaximumAreaAddressesMismatch shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalIn-formation field. The occurrence of this event is counted by the MaximumAreaAddressesMismatches counter.

ownLSPPurge:

Generated when a zero aged copy of a system's own LSP is received from some other system. This represents an erroneous attempt to purge the local system's LSP. The value ISIS.ownLSPPurge shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.communicationsProtocolError. The

perceivedSeverity shall be set to 'Minor'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the ownLSPPurges counter.

originatingLSPBufferSizeMismatch:

Generated when a Level 1 LSP or Level 2 LSP is received which is larger than the local value for originatingL1LSPBufferSize or originatingL2LSPBufferSize respectively or when a Level 1 LSP or Level 2 LSP is received containing the originatingLSPBufferSize option and the value in the PDU option field does not match the local value for originatingL1LSPBufferSize or originatingL2LSPBufferSize respectively. The contents of the LSP header are reported in the additionalInformation field using the notificationLSPHeader parameter. The significance sub-parameter of each item of additionalInformation shall be set to "False" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.originatingLSPBufferSizeMismatch shall be reported in the specificProblems parameter. The probable cause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to "Minor". A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the originatingLSPBufferSizeMismatches counter.

ISPTooLargeToPropagate:

Generated when an attempt is made to propagate an LSP PDU which is larger than the dataLinkBlockSize for a circuit. The contents of the LSP header are reported in the additionalInformation field using the notificationLSPHeader parameter. The significance sub-parameter of each item of additionalInformation shall be set to "False" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.ISPTooLargeToPropagate shall be reported in the specificProblems parameter. The probable cause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to "Major". A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the ISPTooLargeToPropagate counter. ;

cLNSISISBasicImportedInfoNotifications-B BEHAVIOUR DEFINED AS

Imports the communicationsInformation notification from ISO/IEC 10165-5. It is used to report the following protocol events.

sequenceNumberSkip:

Generated when the sequence number of an LSP is incremented by more than one. The value ISIS.sequenceNumberSkip shall be reported in the informationType parameter. No other fields or parameters shall be used, with the exception of parameters in the informationData field. The occurrence of this event is counted by the sequenceNumberSkips counter. ;

ATTRIBUTES

version GET,
iStype INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B GET,
systemId GET,
maximumPathSplits
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.maximumPathSplits-Default
 GET-REPLACE,
minimumLSPTransmissionInterval
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.minimumLSPTransmissionInterval-Default
 GET-REPLACE,
maximumLSPGenerationInterval
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.maximumLSPGenerationInterval-Default
 GET-REPLACE,
minimumBroadcastLSPTransmissionInterval
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.minimumBroadcastLSPTransmissionInterval-Default
 GET-REPLACE,

-- Note: this is defined for all Circuits, but would only be required if one of them were a broadcast Circuit
completeSNPInterval

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.completeSNPInterval-Default

GET-REPLACE,

-- Note: this is defined for all linkages, but would only be required if one of them were a broadcast linkage
originatingL1LSPBufferSize

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.originatingL1LSPBufferSize-Default

GET-REPLACE,

manualAreaAddresses

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.manualAreaAddresses-Default GET ADD-REMOVE,

maximumAreaAddresses

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.maximumAreaAddresses-Default

GET-REPLACE,

minimumLSPGenerationInterval

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.minimumLSPGenerationInterval-Default

GET-REPLACE,

pollESHelloRate

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.pollESHelloRate-Default

GET-REPLACE,

partialSNPInterval

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.partialSNPInterval-Default

GET-REPLACE,

waitingTime

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.waitingTime-Default

GET-REPLACE,

dRISISHelloTimer

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.dRISISHelloTimer-Default

GET-REPLACE,

l1State GET,

areaAddresses GET,

corruptedLSPsDetected GET,

ISPL1DatabaseOverloads GET,

manualAddressesDroppedFromArea GET,

attemptsToExceedMaximumSequenceNumber GET,

sequenceNumberSkips GET,

ownLSPPurges GET,

iDFieldLengthMismatches GET,

maximumAreaAddressesMismatches GET,

originatingLSPBufferSizeMismatches GET,

ISPTooLargeToPropagate GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters

corruptedLSPsDetected

ISPL1DatabaseOverloads

manualAddressesDroppedFromArea

attemptsToExceedMaximumSequenceNumber

sequenceNumberSkips

ownLSPPurges

iDFieldLengthMismatches

maximumAreaAddressesMismatches,

"Rec. X.721 | ISO/IEC 10165-2 : 1992":state


```

    l1State;
NOTIFICATIONS
  "Rec. X.721 | ISO/IEC 10165-2 : 1992":communicationsAlarm
    notificationAreaAddress
    notificationIDLength
    notificationMaximumAreaAddresses
    notificationOverloadStateChange
    notificationReceivingAdjacency
    notificationSourceId
    notificationSystemId
    notificationLSPHeader
    notificationLSPTooLargeToPropagate,
  "Rec. X.723 | ISO/IEC 10165-5":communicationsInformation;
REGISTERED AS {ISIS.poi cLNSISISBasic-P (1)};

```

--11.2.2.2 Level 2 ISIS package attached to cLNS managed object

cLNSISISLevel2-P PACKAGE

BEHAVIOUR cLNSISISLevel2ImportedAlarmNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsAlarm notification from ISO/IEC 10165-2. It is used to report the following protocol events.

ISPL2DatabaseOverload:

generated when the L2State of the system changes between On and Waiting or Waiting and On. The resulting state, shall be reported in the additionalInformation field using the notification-OverloadStateChange parameter and in the case of Waiting the SourceId of the LSP which precipitated the overload shall be reported in the additionalInformation field using the notificationSourceId parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.ISPL2DatabaseOverload shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.resourceAtOrNearing-Capacity. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event shall be counted by the ISPL2Database-Overloads counter.

::

ATTRIBUTES

```

  originatingL2LSPBufferSize
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.originatingL2LSPBufferSize-Default
    GET-REPLACE,

```

l2State GET,

ISPL2DatabaseOverloads GET;

ATTRIBUTE GROUPS

```

  "Rec. X.723 | ISO/IEC 10165-5":counters
    ISPL2DatabaseOverloads,
  "Rec. X.721 | ISO/IEC 10165-2 : 1992":state
    l2State;

```

NOTIFICATIONS

```

  "Rec. X.721 | ISO/IEC 10165-2 : 1992":communicationsAlarm
    notificationOverloadStateChange
    notificationSourceId ;

```

REGISTERED AS {ISIS.poi cLNSISISLevel2-P (2)};

--11.2.2.3 ISIS partition repair package attached to the cLNS managed object

cLNSISISPartitionRepair-P PACKAGE

BEHAVIOUR cLNSISISPartitionRepair-P-ImportedInfoNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsInformation notification from ISO/IEC 10165-5. It is used to report the following protocol events.

partitionVirtualLinkChange:

generated when a virtual link (for the purposes of Level 1 partition repair) is either created or deleted. The direction of the change and the networkEntityTitle of the corresponding virtual adjacency managed object are reported in the communicationData field using the notificationVirtualLinkChange and notificationVirtualLinkAddress parameters respectively. The value ISIS.partitionVirtualLinkChange shall be reported in the informationType parameter. No other fields or parameters shall be used, with the exception of further parameters in the informationData field. The relative order of events relating to the same Virtual Link must be preserved. The occurrence of this event is counted by the partitionVirtualLinkChanges counter.

;;

ATTRIBUTES

maximumVirtualAdjacencies
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.maximumVirtualAdjacencies-Default
 GET-REPLACE,
 partitionAreaAddresses GET,
 partitionDesignatedL2IntermediateSystem GET,
 partitionVirtualLinkChanges GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters
 partitionVirtualLinkChanges;

NOTIFICATIONS

"Rec. X.723 | ISO/IEC 10165-5":communicationsInformation
 notificationVirtualLinkChange
 notificationVirtualLinkAddress;

REGISTERED AS {ISIS.poi cLNSISISPartitionRepair-P (3)};

--11.2.3 ISIS authentication package attached to the cLNS managed object

cLNSISISAuthentication-P PACKAGE

BEHAVIOUR cLNSISISAuthentication-P-ImportedAlarmNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsAlarm notification from ISO/IEC 10165-2. It is used to report the following protocol events.

authenticationFailure:

generated when a PDU is received with an incorrect Authentication information field. The SystemId of the source system shall be reported in the additionalInformation field using the notificationSystemId parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.authenticationFailure shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the authenticationFailures counter.

;;

ATTRIBUTES

areaTransmitPassword
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.password-Default
 GET-REPLACE,
 areaReceivePasswords
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.passwords-Default
 GET-REPLACE
 ADD-REMOVE,
 authenticationFailures GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters
 authenticationFailures;
 NOTIFICATIONS
 "Rec. X.721 | ISO/IEC 10165-2 : 1992":communicationsAlarm
 notificationSystemId;
 REGISTERED AS {ISIS.poi cLNSISISAuthentication-P (4)};

--11.2.3.1 ISIS Level 2 Authentication package attached to the cLNS managed object

cLNSISISLevel2Authentication-P PACKAGE
 ATTRIBUTES
 domainTransmitPassword
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.password-Default
 GET-REPLACE,
 domainReceivePasswords
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.passwords-Default
 GET-REPLACE
 ADD-REMOVE;
 REGISTERED AS {ISIS.poi cLNSISISLevel2Authentication-P (5)};

--11.2.4 Attributes of cLNS object added by conditional packages

areaAddresses ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.AreaAddresses;
 MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR areaAddresses-B BEHAVIOUR
 DEFINED AS
 The union of the sets of manualAreaAddresses reported in all Level 1 Link State PDUs received by this Intermediate System;;
 REGISTERED AS {ISIS.aoi areaAddresses (18)};

areaReceivePasswords ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.Passwords;
 MATCHES FOR EQUALITY, SET-INTERSECTION, SET-COMPARISON;
 BEHAVIOUR areaReceivePasswords-B BEHAVIOUR
 DEFINED AS
 The values to be used as receive passwords to check the receipt of Level 1 LSP, and SNP PDUs;;
 REGISTERED AS {ISIS.aoi areaReceivePasswords (112)};

areaTransmitPassword ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.Password;
 MATCHES FOR EQUALITY;
 BEHAVIOUR areaTransmitPassword-B BEHAVIOUR
 DEFINED AS
 The value to be used as a transmit password in Level 1 LSP, and SNP PDUs transmitted by this Intermediate System;;
 REGISTERED AS {ISIS.aoi areaTransmitPassword (111)};

attemptsToExceedMaximumSequenceNumber ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR attemptsToExceedMaximumSequenceNumber-B BEHAVIOUR
 DEFINED AS
 Number of times the attemptToExceedMaximumSequenceNumber event has been generated;;
 REGISTERED AS {ISIS.aoi attemptsToExceedMaximumSequenceNumber (22)};

completeSNPIInterval ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR completeSNPIInterval-B BEHAVIOUR

DEFINED AS

Interval between generation of Complete Sequence Numbers PDUs by a Designated Intermediate System on a broadcast circuit;;

resettingTimer-B;

REGISTERED AS {ISIS.aoi completeSNPInterval (8)};

corruptedLSPsDetected ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR corruptedLSPsDetected-B BEHAVIOUR

DEFINED AS Number of Corrupted LSP Detected events generated;;

REGISTERED AS {ISIS.aoi corruptedLSPsDetected (19)};

domainTransmitPassword ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Password;

MATCHES FOR EQUALITY;

BEHAVIOUR domainTransmitPassword-B BEHAVIOUR

DEFINED AS

The value to be used as a transmit password in Level 2 LSP, and SNP PDUs transmitted by this Intermediate System;;

REGISTERED AS {ISIS.aoi domainTransmitPassword (113)};

domainReceivePasswords ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Passwords;

MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;

BEHAVIOUR domainReceivePasswords-B BEHAVIOUR

DEFINED AS

The values to be used as receive passwords to check the receipt of Level 2 LSP, and SNP PDUs;;

REGISTERED AS {ISIS.aoi domainReceivePasswords (114)};

dRISISHelloTimer ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR dRISISHelloTimer-B BEHAVIOUR

DEFINED AS

The interval between the generation of IIH PDUs by the designated IS on a LAN;;

resettingHoldingTimer-B;

REGISTERED AS {ISIS.aoi dRISISHelloTimer (16)};

iDFieldLengthMismatches ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR iDFieldLengthMismatches-B BEHAVIOUR

DEFINED AS Number of times the iDFieldLengthMismatch event has been generated;;

REGISTERED AS {ISIS.aoi iDFieldLengthMismatches (25)};

iSType ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.ISType;

MATCHES FOR EQUALITY;

BEHAVIOUR iSType-B BEHAVIOUR

DEFINED AS

The type of this Intermediate System. The value of this attribute is only settable via the create parameter;;

REGISTERED AS {ISIS.aoi iSType (2)};

l1State ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.DatabaseState;

MATCHES FOR EQUALITY;

BEHAVIOUR l1State-B BEHAVIOUR

DEFINED AS The state of the Level 1 database;;

REGISTERED AS {ISIS.aoi l1State (17)};

l2State ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.DatabaseState;

MATCHES FOR EQUALITY;
 BEHAVIOUR l2State-B BEHAVIOUR
 DEFINED AS The state of the Level 2 database;;
 REGISTERED AS {ISIS.aoi l2State (28)};

ISPL1DatabaseOverloads ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR ISPL1DatabaseOverloads-B BEHAVIOUR
 DEFINED AS Number of times the l1LSPDatabaseOverload event has been generated;;
 REGISTERED AS {ISIS.aoi ISPL1DatabaseOverloads (20)};

ISPL2DatabaseOverloads ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR ISPL2DatabaseOverloads-B BEHAVIOUR
 DEFINED AS Number of times the l2LSPDatabaseOverload event has been generated;;
 REGISTERED AS {ISIS.aoi ISPL2DatabaseOverloads (32)};

manualAddressesDroppedFromArea ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR manualAddressesDroppedFromArea-B BEHAVIOUR
 DEFINED AS Number of times the manualAddressesDroppedFromArea event has been generated;;
 REGISTERED AS {ISIS.aoi manualAddressesDroppedFromArea (21)};

manualAreaAddresses ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.AreaAddresses;
 MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR manualAreaAddresses-B BEHAVIOUR
 DEFINED AS
 Area Addresses to be used for this Intermediate System. At least one value must be supplied. The maximum number of Area Addresses which may exist in the set is MaximumAreaAddresses;;
 REGISTERED AS {ISIS.aoi manualAreaAddresses (10)};

maximumAreaAddresses ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.MaximumAreaAddresses;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR maximumAreaAddresses-B BEHAVIOUR
 DEFINED AS
 maximum number of area addresses to be permitted for this ISs area. Note that all ISs in the area must have the same value configured for this attribute if correct operation is to be assured.,,
 replaceOnlyWhileDisabled-B;
 PARAMETERS constraintViolation;
 REGISTERED AS {ISIS.aoi maximumAreaAddresses (4)};

maximumAreaAddressesMismatches ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR maximumAreaAddressesMismatches-B BEHAVIOUR
 DEFINED AS Count of maximumAreaAddressesMismatch events.;;
 REGISTERED AS {ISIS.aoi maximumAreaAddressesMismatches(118)};

maximumLSPGenerationInterval ATTRIBUTE
 DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR maximumLSPGenerationInterval-B BEHAVIOUR
 DEFINED AS Maximum interval between generated LSPs by this system,,
 resettingTimer-B;
 REGISTERED AS {ISIS.aoi maximumLSPGenerationInterval (6)};

maximumPathSplits ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.MaximumPathSplits;
 MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR maximumPathSplits-B BEHAVIOUR

DEFINED AS Maximum number of paths which it is permitted to split traffic between;
replaceOnlyWhileDisabled-B;

PARAMETERS constraintViolation;

REGISTERED AS {ISIS.aoi maximumPathSplits (3)};

maximumVirtualAdjacencies ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.MaximumVirtualAdjacencies;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR maximumVirtualAdjacencies-B BEHAVIOUR

DEFINED AS Maximum number of Virtual Adjacencies which may be created to repair partitioned Level 1 domains;
resourceLimiting-B;

PARAMETERS constraintViolation;

REGISTERED AS {ISIS.aoi maximumVirtualAdjacencies (27)};

minimumBroadcastLSPTransmissionInterval ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR minimumBroadcastLSPTransmissionInterval-B BEHAVIOUR

DEFINED AS

Minimum interval between transmission of LSPs on a broadcast circuit (See clause 7.3.15.6). This timer shall be capable of a resolution not coarser than 10 ms;;

resettingTimer-B;

REGISTERED AS {ISIS.aoi minimumBroadcastLSPTransmissionInterval (7)};

minimumLSPGenerationInterval ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR minimumLSPGenerationInterval-B BEHAVIOUR

DEFINED AS Minimum interval between successive generation of LSPs with the same LSPID by this IS;;
resettingTimer-B;

REGISTERED AS {ISIS.aoi minimumLSPGenerationInterval (11)};

minimumLSPTransmissionInterval ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR minimumLSPTransmissionInterval-B BEHAVIOUR

DEFINED AS Minimum interval between re- transmissions of an LSP;;
resettingTimer-B;

REGISTERED AS {ISIS.aoi minimumLSPTransmissionInterval (5)};

originatingL1LSPBufferSize ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.OriginatingLSPBufferSize;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR originatingL1LSPBufferSize-B BEHAVIOUR

DEFINED AS The maximum size of Level 1 LSPs and SNPs originated by this system;;
replaceOnlyWhileDisabled-B;

PARAMETERS constraintViolation;

REGISTERED AS {ISIS.aoi originatingL1LSPBufferSize (9)};

originatingL2LSPBufferSize ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.OriginatingLSPBufferSize;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR originatingL2LSPBufferSize-B BEHAVIOUR

DEFINED AS The maximum size of Level 2 LSPs and SNPs originated by this system;;
replaceOnlyWhileDisabled-B;

PARAMETERS constraintViolation;

REGISTERED AS {ISIS.aoi originatingL2LSPBufferSize (26)};

ownLSPPurges ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR ownLSPPurges-B BEHAVIOUR
 DEFINED AS Number of times the ownLSPPurged event has been generated;;
 REGISTERED AS {ISIS.aoi ownLSPPurges (24)};

partialSNPInterval ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR partialSNPInterval-B BEHAVIOUR
 DEFINED AS Minimum interval between sending Partial Sequence Number PDUs,
 resettingTimer-B;
 REGISTERED AS {ISIS.aoi partialSNPInterval (14)};

partitionAreaAddresses ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.AreaAddresses;
 MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR partitionAreaAddresses-B BEHAVIOUR
 DEFINED AS
 The set union of all manualAreaAddresses of all Intermediate systems in the partition reachable by non-virtual links
 (calculated from their Level 1 LSPs);
 REGISTERED AS {ISIS.aoi partitionAreaAddresses (29)};

partitionDesignatedL2IntermediateSystem ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SystemId;
 MATCHES FOR EQUALITY;
 BEHAVIOUR partitionDesignatedL2IntermediateSystem-B BEHAVIOUR
 DEFINED AS The ID of the Partition Designated Level 2 Intermediate System for this system;;
 REGISTERED AS {ISIS.aoi partitionDesignatedL2IntermediateSystem (30)};

partitionVirtualLinkChanges ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR partitionVirtualLinkChanges-B BEHAVIOUR
 DEFINED AS Number of times the partitionVirtualLink change event has been generated;;
 REGISTERED AS {ISIS.aoi partitionVirtualLinkChanges (31)};

pollESHelloRate ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR pollESHelloRate-B BEHAVIOUR
 DEFINED AS The value to be used for the suggested ES configuration timer in ISH PDUs when soliciting the ES
 configuration;;
 REGISTERED AS {ISIS.aoi pollESHelloRate (13)};

sequenceNumberSkips ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR sequenceNumberSkips-B BEHAVIOUR
 DEFINED AS Number of times the sequenceNumberSkipped event has been generated;;
 REGISTERED AS {ISIS.aoi sequenceNumberSkips (23)};

systemId ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SystemId;
 MATCHES FOR EQUALITY;
 BEHAVIOUR systemId-B BEHAVIOUR
 DEFINED AS
 The ID for the local system, to be appended to each of the system's area address(es) to form the Network Entity Titles
 valid for this IS. The derivation of a value for SystemId is a local matter;;
 REGISTERED AS {ISIS.aoi systemId (119)};

version ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Version;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR version-B BEHAVIOUR

DEFINED AS The edition of this International Standard to which the implementation conforms;;

REGISTERED AS {ISIS.aoi version (1)};

waitingTime ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR waitingTime-B BEHAVIOUR

DEFINED AS Amount of time to delay in waiting state before entering On state;;

resettingTimer-B;

REGISTERED AS {ISIS.aoi waitingTime (15)};

originatingLSPBufferSizeMismatches ATTRIBUTE

DERIVED FROM "Rec. X. 723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR originatingLSPBufferSizeMismatches-B BEHAVIOUR

DEFINED AS Count of originatingLSPBufferSizeMismatch events;;

REGISTERED AS {ISIS.aoi originatingLSPBufferSizeMismatches(124)};

ISPTooLargeToPropagate ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR ISPTooLargeToPropagate-B BEHAVIOUR

DEFINED AS Count of ISPTooLargeToPropagate events;;

REGISTERED AS {ISIS.aoi ISPTooLargeToPropagate(125)};

--11.2.5 Conditional packages for linkage MO defined in ISO 10733

--11.2.5.1 Basic ISIS linkage package

linkageISISBasic-P PACKAGE

BEHAVIOUR linkageISISBasicImportedAlarmNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsAlarm notification from ISO/IEC 10165-2. It is used to report the following protocol events.

versionSkew:

generated when an attempt to initialise with an adjacent system fails as a result of the versions of the protocol are not compatible. The protocol version field from the received PDU shall be reported in the additionalInformation field using the notificationVersion parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.versionSkew shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.versionMismatch. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the initialisationFailures counter.

areaMismatch:

generated when an attempt to initialise with an adjacent system fails as a result of two level 1 ISs not sharing any area addresses in common. The AREA ADDRESSES field and the ID field of the received PDU are reported in the additionalInformation field using the notificationAreaAddresses and notificationSystemId parameters respectively. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.areaMismatch shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the initialisationFailures counter.

rejectedAdjacency:

generated when an attempt to create a new adjacency is rejected, because of a lack of resources. The ID field of the received PDU is reported in the additionaldata field using the notificationSystemId parameter. The value ISIS.rejectedAdjacency shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.resourceAtOrNearingCapacity. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the rejectedAdjacencies counter.

iDFieldLengthMismatch:

generated when a PDU is received with a different value for ID field length to that of the receiving Intermediate system. The received ID LENGTH and SOURCE ID are reported in the additionalInformation field using the notificationIDLength and notificationSourceId parameters respectively. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.iDFieldLengthMismatch shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the iDFieldLengthMismatches counter.

maximumAreaAddressesMismatch:

generated when a PDU is received with a different value for maximumAreaAddresses from that of the receiving Intermediate system. The received MAXIMUM AREA ADDRESSES and SOURCE ID are reported in the additionalInformation field using the notification maximumAreaAddresses and notificationSourceId parameters respectively. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.maximumAreaAddressesMismatch shall be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the maximumAreaAddressesMismatches counter.

;;

ATTRIBUTES

type INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B GET,

iSISHelloTimer

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.iSISHelloTimer-Default
 GET-REPLACE,

11DefaultMetric

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.defaultMetric-Default
 PERMITTED VALUES ISIS.DefaultMetric-Permitted
 GET-REPLACE,

11DelayMetric

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,

11ExpenseMetric

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,

11ErrorMetric

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,

externalDomain

REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.externalDomain-Default
 GET-REPLACE,

changesInAdjacencyState GET,
initialisationFailures GET,
rejectedAdjacencies GET,
iSISControlPDUsSent GET,
iSISControlPDUsReceived GET,
iDFieldLengthMismatches GET,
maximumAreaAddressesMismatches GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters
changesInAdjacencyState
initialisationFailures
rejectedAdjacencies
iSISControlPDUsSent
iSISControlPDUsReceived
iDFieldLengthMismatches;

NOTIFICATIONS

"Rec. X.721 | ISO/IEC 10165-2 : 1992":communicationsAlarm
notificationAreaAddresses
notificationIDLength
notificationMaximumAreaAddresses
notificationSourceId
notificationSystemId
notificationVersion;

REGISTERED AS {ISIS.poi linkageISISBasic-P (6)};

--11.2.5.2 ISIS broadcast linkage package

linkageISISBroadcast-P PACKAGE

BEHAVIOUR linkageISISBroadcastImportedInfoNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsInformation notification from ISO/IEC 10165-5. It is used to report the following protocol events.

lanL1DesignatedIntermediateSystemChange:

generated when the local system either elects itself or resigns as being the LAN L1 Designated Intermediate System on this circuit. The direction of the change is reported in the communicationData field using the notificationDesignated-IntermediateSystemChange parameter. The value ISIS.lanL1DesignatedIntermediateSystemChange shall be reported in the informationType parameter. No other fields or parameters shall be used, with the exception of further parameters in the informationData field. The relative order of these events must be preserved. The occurrence of this event is counted by the lanL1DesignatedIntermediateSystemChanges counter.

;;

ATTRIBUTES

l1IntermediateSystemPriority
REPLACE-WITH-DEFAULT
DEFAULT VALUE ISIS.l1IntermediateSystemPriority-Default
GET-REPLACE,

l1CircuitID GET,
l1DesignatedIntermediateSystem GET,
lanL1DesignatedIntermediateSystemChanges GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters
lanL1DesignatedIntermediateSystemChanges;

NOTIFICATIONS

"Rec. X.723 | ISO/IEC 10165-5":communicationsInformation
notificationDesignatedIntermediateSystemChange;

REGISTERED AS {ISIS.poi linkageISISBroadcast-P (7)};

--11.2.5.3 ISIS point-to-point linkage package

linkageISISPtToPt-P PACKAGE
 ATTRIBUTES
 ptPtCircuitID GET;
 REGISTERED AS {ISIS.poi linkageISISPtToPt-P (8)};

--11.2.5.4 ISIS call establishment metric increment linkage package

linkageISISDASCallEstablishmentMetricIncrement-P PACKAGE
 BEHAVIOUR linkageISISDASCallEstablishmentMetricIncrement-P-B BEHAVIOUR
 DEFINED AS
 Present when values of call establishment metric increment greater than zero are supported and the parent cLNS MO has iSType Level2;;
 ATTRIBUTES
 callEstablishmentDefaultMetricIncrement
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.callEstablishmentMetricIncrement-Default
 GET-REPLACE,
 callEstablishmentDelayMetricIncrement
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.callEstablishmentMetricIncrement-Default
 GET-REPLACE,
 callEstablishmentExpenseMetricIncrement
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.callEstablishmentMetricIncrement-Default
 GET-REPLACE,
 callEstablishmentErrorMetricIncrement
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.callEstablishmentMetricIncrement-Default
 GET-REPLACE;
 REGISTERED AS {ISIS.poi linkageISISDASCallEstablishmentMetricIncrement-P (9)};

--11.2.5.5 ISIS static linkage package

linkageISISStatic-P PACKAGE
 ATTRIBUTES
 outgoingCallIVMO
 REPLACE-WITH-DEFAULT
 GET-REPLACE,
 ptPtCircuitID GET,
 neighbourSNPAAddress
 REPLACE-WITH-DEFAULT
 GET-REPLACE;
 REGISTERED AS {ISIS.poi linkageISISStatic-P (11)};

--11.2.5.6 ISIS level 2 linkage package

linkageISISLevel2-P PACKAGE
 ATTRIBUTES
 l2DefaultMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.defaultMetric-Default
 PERMITTED VALUES ISIS.DefaultMetric-Permitted
 GET-REPLACE,
 l2DelayMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,

```

l2ExpenseMetric
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.optionalMetric-Default
    GET-REPLACE,
l2ErrorMetric
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.optionalMetric-Default
    GET-REPLACE,
manualL2OnlyMode
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.manualL2OnlyMode-Default
    GET-REPLACE;
REGISTERED AS {ISIS.poi linkageISISLevel2-P (13)};

```

--11.2.5.7 ISIS level 2 broadcast linkage package

iSISLevel2Broadcast-P PACKAGE

BEHAVIOUR linkageISISLevel2BroadcastImportedInfoNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsInformation notification from ISO/IEC 10165-5. It is used to report the following protocol events.

lanL2DesignatedIntermediateSystemChange:

generated when the local system either elects itself or resigns as being the LAN L2 Designated Intermediate System on this circuit. The direction of the change is reported in the communicationData field using the notificationDesignatedIntermediate-SystemChange parameter. The value ISIS.lanL2DesignatedIntermediateSystemChange shall be reported in the informationType parameter. No other fields or parameters shall be used, with the exception of further parameters in the informationData field. The relative order of these events must be preserved. The occurrence of this event is counted by the lanL2DesignatedIntermediate-SystemChanges counter.

::

ATTRIBUTES

```

l2IntermediateSystemPriority
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.l2IntermediateSystemPriority-Default
    GET-REPLACE,

```

```

l2CircuitID GET,

```

```

l2DesignatedIntermediateSystem GET,

```

```

lanL2DesignatedIntermediateSystemChanges GET;

```

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters

lanL2DesignatedIntermediateSystemChanges;

NOTIFICATIONS

"Rec. X.723 | ISO/IEC 10165-5":communicationsInformation
notificationDesignatedIntermediateSystemChange;

REGISTERED AS {ISIS.poi ISISLevel2Broadcast-P (14)};

--11.2.5.8 ISIS linkage authentication package

linkageISISAuthentication-P PACKAGE

BEHAVIOUR linkageISISAuthentication-P-ImportedAlarmNotifications-B BEHAVIOUR

DEFINED AS

Imports the communicationsAlarm notification from ISO/IEC 10165-2. It is used to report the following protocol events.

authenticationFailure:

generated when a PDU is received with an incorrect Authentication information field. The SystemId of the source system is reported in the additionalInformation field using the notificationSystemId parameter. The significance sub-parameter of each item of additionalInformation shall be set to the value "false" (i.e. not significant) so that a managing system receiving the event report will be less likely to reject it. The value ISIS.authenticationFailure shall

be reported in the specificProblems parameter. The probableCause shall be set to NLM.configurationOrCustomisationError. The perceivedSeverity shall be set to 'Major'. A subsequent communicationsAlarm with a perceivedSeverity value of "Cleared" shall not be generated. No other fields or parameters shall be used, with the exception of further parameters in the additionalInformation field. The occurrence of this event is counted by the authenticationFailures counter.

;;

ATTRIBUTES

circuitTransmitPassword

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.password-Default

GET-REPLACE,

circuitReceivePasswords

REPLACE-WITH-DEFAULT

DEFAULT VALUE ISIS.passwords-Default

GET-REPLACE ADD-REMOVE,

authenticationFailures GET;

ATTRIBUTE GROUPS

"Rec. X.723 | ISO/IEC 10165-5":counters

authenticationFailures;

NOTIFICATIONS

"Rec. X.721 | ISO/IEC 10165-2 : 1992":communicationsAlarm

notificationSystemId;

REGISTERED AS {ISIS.poi linkageISISAuthentication-P (15)};

--11.2.5.9 Attributes for the linkage MO from ISO 10733 added by ISIS conditional packages

callEstablishmentDefaultMetricIncrement ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR callEstablishmentDefaultMetricIncrement-B BEHAVIOUR

DEFINED AS Additional value to be reported for the default metric value of unestablished DA adjacencies;;

REGISTERED AS {ISIS.aoi callEstablishmentDefaultMetricIncrement (52)};

callEstablishmentDelayMetricIncrement ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR callEstablishmentDelayMetricIncrement-B BEHAVIOUR

DEFINED AS Additional value to be reported for the delay metric value of unestablished DA adjacencies;;

REGISTERED AS {ISIS.aoi callEstablishmentDelayMetricIncrement (53)};

callEstablishmentErrorMetricIncrement ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR callEstablishmentErrorMetricIncrement-B BEHAVIOUR

DEFINED AS Additional value to be reported for the Error metric value of unestablished DA adjacencies;;

REGISTERED AS {ISIS.aoi callEstablishmentErrorMetricIncrement (55)};

callEstablishmentExpenseMetricIncrement ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR callEstablishmentExpenseMetricIncrement-B BEHAVIOUR

DEFINED AS Additional value to be reported for the Expense metric value of unestablished DA adjacencies;;

REGISTERED AS {ISIS.aoi callEstablishmentExpenseMetricIncrement (54)};

changesInAdjacencyState ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR changesInAdjacencyState-B BEHAVIOUR

DEFINED AS Number of Adjacency State Change events generated;;

REGISTERED AS {ISIS.aoi changesInAdjacencyState (40)};

circuitReceivePasswords ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Passwords;
 MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR circuitReceivePasswords-B BEHAVIOUR
 DEFINED AS The values to be used as receive passwords to check the receipt of IIH PDUs;;
 REGISTERED AS {ISIS.aoi circuitReceivePasswords (116)};

circuitTransmitPassword ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Password;
 MATCHES FOR EQUALITY;
 BEHAVIOUR circuitTransmitPassword-B BEHAVIOUR
 DEFINED AS The value to be used as a transmit password in IIH PDUs transmitted by this Intermediate System;;
 REGISTERED AS {ISIS.aoi circuitTransmitPassword (115)};

externalDomain ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Boolean;
 MATCHES FOR EQUALITY;
 BEHAVIOUR externalDomain-B BEHAVIOUR
 DEFINED AS If TRUE, suppress normal transmission of and interpretation of Intra-domain ISIS PDUs on this circuit;;
 REGISTERED AS {ISIS.aoi externalDomain (46)};

initialisationFailures ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR initialisationFailures-B BEHAVIOUR
 DEFINED AS Number of Initialisation Failure events generated;;
 REGISTERED AS {ISIS.aoi initialisationFailures (41)};

iSISControlPDUsReceived ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR iSISControlPDUsReceived-B BEHAVIOUR
 DEFINED AS Number of control PDUs received on this circuit;;
 REGISTERED AS {ISIS.aoi iSISControlPDUsReceived (44)};

iSISControlPDUsSent ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
 BEHAVIOUR iSISControlPDUsSent-B BEHAVIOUR
 DEFINED AS Number of control PDUs sent on this circuit;;
 REGISTERED AS {ISIS.aoi iSISControlPDUsSent (43)};

iSISHelloTimer ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR iSISHelloTimer-B BEHAVIOUR
 DEFINED AS The period between IIH PDUs. It is also used as the period between ISH PDUs when polling the ES configuration;;
 resettingHoldingTimer-B;
 REGISTERED AS {ISIS.aoi iSISHelloTimer (45)};

l1CircuitID ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.CircuitID;
 MATCHES FOR EQUALITY;
 BEHAVIOUR l1CircuitID-B BEHAVIOUR
 DEFINED AS
 The LAN ID allocated by the LAN Level 1 Designated Intermediate System. Where this system is not aware of the value (because it is not participating in the Level 1 Designated Intermediate System election), this attribute has the value which would be proposed for this circuit. (i.e. the concatenation of the local system ID and the one octet local Circuit ID for this circuit;;
 REGISTERED AS {ISIS.aoi l1CircuitID (48)};

11ErrorMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR 11ErrorMetric-B BEHAVIOUR
 DEFINED AS

The error metric value of this circuit for Level 1 traffic. The value of zero is reserved to indicate that this metric is not supported;;

REGISTERED AS {ISIS.aoi 11ErrorMetric (38)};

11ExpenseMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR 11ExpenseMetric-B BEHAVIOUR
 DEFINED AS

The expense metric value of this circuit for Level 1 traffic. The value of zero is reserved to indicate that this metric is not supported;;

REGISTERED AS {ISIS.aoi 11ExpenseMetric (37)};

11DefaultMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR 11defaultMetric-B BEHAVIOUR

DEFINED AS The default metric value of this circuit for Level 1 traffic. ;;

REGISTERED AS {ISIS.aoi 11DefaultMetric (35)};

11DelayMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR 11DelayMetric-B BEHAVIOUR
 DEFINED AS

The delay metric value of this circuit for Level 1 traffic. The value of zero is reserved to indicate that this metric is not supported;;

REGISTERED AS {ISIS.aoi 11DelayMetric (36)};

11DesignatedIntermediateSystem ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SystemId;
 MATCHES FOR EQUALITY;
 BEHAVIOUR 11DesignatedIntermediateSystem-B BEHAVIOUR
 DEFINED AS

The ID of the LAN Level 1 Designated Intermediate System on this circuit. If, for any reason this system is not partaking in the relevant Designated Intermediate System election process, then the value returned is the zero length OCTET STRING;;

REGISTERED AS {ISIS.aoi 11DesignatedIntermediateSystem (49)};

11IntermediateSystemPriority ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.IntermediateSystemPriority;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR 11IntermediateSystemPriority-B BEHAVIOUR

DEFINED AS Priority for becoming LAN Level 1 Designated Intermediate System;;

REGISTERED AS {ISIS.aoi 11IntermediateSystemPriority (47)};

12CircuitID ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.CircuitID;
 MATCHES FOR EQUALITY;
 BEHAVIOUR 12CircuitID-B BEHAVIOUR
 DEFINED AS

The LAN ID allocated by the LAN Level 2 Designated Intermediate System. Where this system is not aware of the value (because it is not participating in the Level 2 Designated Intermediate System election), this attribute has the value which would be proposed for this circuit. (i.e. the concatenation of the local system ID and the one octet local Circuit ID for this circuit.;;

REGISTERED AS {ISIS.aoi 12CircuitID (74)};

l2DefaultMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR l2defaultMetric-B BEHAVIOUR
DEFINED AS The default metric value of this circuit for Level 2 traffic. ;
REGISTERED AS {ISIS.aoi l2DefaultMetric (68)};

l2DelayMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR l2DelayMetric-B BEHAVIOUR
DEFINED AS
The delay metric value of this circuit for Level 2 traffic. The value of zero is reserved to indicate that this metric is not supported;;
REGISTERED AS {ISIS.aoi l2DelayMetric (69)};

l2DesignatedIntermediateSystem ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SystemId;
MATCHES FOR EQUALITY;
BEHAVIOUR l2DesignatedIntermediateSystem-B BEHAVIOUR
DEFINED AS
The ID of the LAN Level 2 Designated Intermediate System on this circuit. If, for any reason this system is not partaking in the relevant Designated Intermediate System election process, then the value returned is the zero length OCTET STRING;;
REGISTERED AS {ISIS.aoi l2DesignatedIntermediateSystem (75)};

l2ErrorMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR l2ErrorMetric-B BEHAVIOUR
DEFINED AS
The error metric value of this circuit for Level 2 traffic. The value of zero is reserved to indicate that this metric is not supported;;
REGISTERED AS {ISIS.aoi l2ErrorMetric (71)};

l2ExpenseMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR l2ExpenseMetric-B BEHAVIOUR
DEFINED AS
The expense metric value of this circuit for Level 2 traffic. The value of zero is reserved to indicate that this metric is not supported;;
REGISTERED AS {ISIS.aoi l2ExpenseMetric (70)};

l2IntermediateSystemPriority ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.IntermediateSystemPriority;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR l2IntermediateSystemPriority-B BEHAVIOUR
DEFINED AS Priority for becoming LAN Level 2 Designated Intermediate System;;
REGISTERED AS {ISIS.aoi l2IntermediateSystemPriority (73)};

lanL1DesignatedIntermediateSystemChanges ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
BEHAVIOUR lanL1DesignatedIntermediateSystemChanges-B BEHAVIOUR
DEFINED AS Number of LAN L1 Designated Intermediate System Change events generated;;
REGISTERED AS {ISIS.aoi lanL1DesignatedIntermediateSystemChanges (50)};

lanL2DesignatedIntermediateSystemChanges ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;
BEHAVIOUR lanL2DesignatedIntermediateSystemChanges-B BEHAVIOUR

DEFINED AS Number of LAN L2 Designated Intermediate System Change events generated;;
 REGISTERED AS {ISIS.aoi lanL2DesignatedIntermediateSystemChanges (76)};

manualL2OnlyMode ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.Boolean;

MATCHES FOR EQUALITY;

BEHAVIOUR manualL2OnlyMode-B BEHAVIOUR

DEFINED AS

When True, indicates that this Circuit is to be used only for Level 2;; replaceOnlyWhileDisabled-B; PARAMETERS
 constraintViolation;

REGISTERED AS {ISIS.aoi manualL2OnlyMode (72)};

outgoingCallIVMO ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.LocalDistinguishedName;

MATCHES FOR EQUALITY;

BEHAVIOUR outgoingCallIVMO-B BEHAVIOUR

DEFINED AS

reference to the virtualCallIVMO to be used to establish communication with a neighbour over this circuit. This
 IVMO contains, among other things, the SNPA Address to use as the called address.;;

REGISTERED AS {ISIS.aoi outgoingCallIVMO (120)};

ptPtCircuitID ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.CircuitID;

MATCHES FOR EQUALITY;

BEHAVIOUR ptPtCircuitID-B BEHAVIOUR

DEFINED AS

The ID of the circuit allocated during initialisation. If no value has been negotiated (either because the adjacency is to
 an End system, or because initialisation has not yet successfully completed), this attribute has the value which would
 be proposed for this circuit. (i.e. the concatenation of the local system ID and the one octet local Circuit ID for this
 circuit.;;

REGISTERED AS {ISIS.aoi ptPtCircuitID (51)};

rejectedAdjacencies ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":nonWrappingCounter;

BEHAVIOUR rejectedAdjacencies-B BEHAVIOUR

DEFINED AS Number of Rejected Adjacency events generated;;

REGISTERED AS {ISIS.aoi rejectedAdjacencies (42)};

type ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.CircuitType;

MATCHES FOR EQUALITY;

BEHAVIOUR type-B BEHAVIOUR

DEFINED AS The type of the circuit. This attribute may only be set when the linkage MO is created. Subsequently it is
 read-only;;

REGISTERED AS {ISIS.aoi type (33)};

--11.2.6 Adjacency managed object

- Created either through the adjacency-linkage Name binding for
- adjacencies instantiated by protocol operation, or the
- adjacency-linkage-management name binding for adjacencies created
- via explicit system management operation --

adjacency MANAGED OBJECT CLASS

DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992":top;

CHARACTERIZED BY adjacency-P PACKAGE

BEHAVIOUR adjacencyStateChange-B BEHAVIOUR

DEFINED AS

This Managed Object imports the 10165-2 stateChange notification. It is used to report changes to the adjacencyState
 attribute. The value "Down" shall be reported when the adjacency is deleted. A single parameter set shall be included

in the State Change definition field. Only the (mandatory) attributeId and newAttributeValue parameters shall be used.

--

ATTRIBUTES

adjacencyId GET,
adjacencyState GET, -- Note: this is NOT operational state
neighbourSNPAAAddress
INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B
GET,
neighbourSystemType GET,
neighbourSystemIds
INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B
GET;

NOTIFICATIONS

"Rec. X.721 | ISO/IEC 10165-2 : 1992":stateChange; ;

CONDITIONAL PACKAGES

iSAdjacency-P PRESENT IF

"the adjacency is to an IS (i.e the neighbourSystemType is Intermediate System, L1 Intermediate System or L2 Intermediate System)", broadcastISAdjacency-P PRESENT IF the parent Linkage MO is of type broadcast and is to an IS as above;

REGISTERED AS {ISIS.moi adjacency (1)};

--11.2.6.1 IS adjacency package of the adjacency managed object

iSAdjacency-P PACKAGE

ATTRIBUTES

adjacencyUsage GET,
areaAddressesOfNeighbour GET,
holdingTimer GET;

REGISTERED AS {ISIS.poi iSAdjacency-P (19)};

--11.2.6.2 Broadcast IS adjacency package of the adjacency managed object

broadcastISAdjacency-P PACKAGE

ATTRIBUTES

priorityOfNeighbour GET;

REGISTERED AS {ISIS.poi broadcastISAdjacency-P (20)};

--11.2.6.3 Name bindings for the adjacency managed object

adjacency-linkage NAME BINDING

SUBORDINATE OBJECT CLASS adjacency AND SUBCLASSES;

NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":linkage AND SUBCLASSES;

WITH ATTRIBUTE adjacencyId;

BEHAVIOUR adjacency-linkage-B BEHAVIOUR

DEFINED AS

This Name Binding is used for adjacencies created automatically by operation of the ISO/IEC 10589 protocol machine.;

adjacencyId-B;

REGISTERED AS {ISIS.nboi adjacency-linkage (7)};

adjacency-linkage-management NAME BINDING

SUBORDINATE OBJECT CLASS adjacency AND SUBCLASSES;

NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":linkage AND SUBCLASSES;

WITH ATTRIBUTE adjacencyId;

BEHAVIOUR adjacency-linkage-management-B BEHAVIOUR

DEFINED AS This Name Binding is used for adjacencies created by system management.;;
 adjacencyId-B;
 CREATE WITH-REFERENCE-OBJECT reservedName;
 DELETE ONLY-IF-NO-CONTAINED-OBJECTS;
 REGISTERED AS {ISIS.nboi adjacency-linkage-management (8)};;

--11.2.6.4 Attributes of the adjacency managed object

adjacencyId ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.GraphicString;
 MATCHES FOR EQUALITY, SUBSTRINGS;
 BEHAVIOUR adjacencyId-B BEHAVIOUR
 DEFINED AS

A string which is the Identifier for the Adjacency and which is unique amongst the set of Adjacencies main-tained for this linkage. If this is an adjacency created by system management, it is set by the System Manager when the Adjacency is created, otherwise it is generated by the implementation such that it is unique. The set of identifiers containing the leading string "Auto" are reserved for Automatic Adjacencies. An attempt by system management to create an adjacency with such an identifier will cause a reserved name violation;;

REGISTERED AS {ISIS.aoi adjacencyId (77)};;

adjacencyState ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.AdjacencyState;
 MATCHES FOR EQUALITY;
 BEHAVIOUR adjacencyState-B BEHAVIOUR
 DEFINED AS The state of the adjacency;;
 REGISTERED AS {ISIS.aoi adjacencyState (78)};;

adjacencyUsage ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.AdjacencyUsage;
 MATCHES FOR EQUALITY;
 BEHAVIOUR adjacencyUsage-B BEHAVIOUR
 DEFINED AS

The usage of the Adjacency. An Adjacency of type Level 1 will be used for Level 1 traffic only. An adjacency of type Level 2 will be used for Level 2 traffic only. An adjacency of type Level 1 and 2 will be used for both Level 1 and Level 2 traffic. There may be two adjacencies (of types Level 1 and Level 2 between the same pair of Intermediate Systems.;;

REGISTERED AS {ISIS.aoi adjacencyUsage (82)};;

areaAddressesOfNeighbour ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.AreaAddresses;
 MATCHES FOR EQUALITY SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR areaAddressesOfNeighbour-B BEHAVIOUR
 DEFINED AS This contains the Area Addresses of a neighbour Intermediate System from the IIH PDU.;;
 REGISTERED AS {ISIS.aoi areaAddressesOfNeighbour (84)};;

holdingTimer ATTRIBUTE

DERIVED FROM "Rec. X.723 | ISO/IEC 10165-5":timer;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR holdingTimer-B BEHAVIOUR
 DEFINED AS Holding time for this adjacency updated from the IIH PDUs;;
 REGISTERED AS {ISIS.aoi holdingTimer (85)};;

priorityOfNeighbour ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.IntermediateSystemPriority;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR priorityOfNeighbour-B BEHAVIOUR
 DEFINED AS

Priority of neighbour on this adjacency for becoming LAN Level 1 Designated Intermediate System if adjacencyType is L1 Intermediate System or LAN Level 2 Designated Intermediate System if adjacencyType is L2 Intermediate System;;

REGISTERED AS {ISIS.aoi priorityOfNeighbour (86)};;

neighbourSNPAAddress ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SNPAAddress;
MATCHES FOR EQUALITY;
BEHAVIOUR neighbourSNPAAddress-B BEHAVIOUR
DEFINED AS The SNPA address of the neighbour system;
replaceOnlyWhileDisabled-B;
PARAMETERS constraintViolation;
REGISTERED AS {ISIS.aoi neighbourSNPAAddress (79)};

neighbourSystemIds ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SystemIds;
MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
BEHAVIOUR neighbourSystemIds-B BEHAVIOUR
DEFINED AS
For Intermediate system neighbours: Contains the single SystemId of the neighbouring Intermediate system obtained from the Source ID field of the neighbour's IIH PDU. For End system neighbours: Contains the set of system ID(s) of a neighbour End system.;;
REGISTERED AS {ISIS.aoi neighbourSystemIds (83)};

neighbourSystemType ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.NeighbourSystemType;
MATCHES FOR EQUALITY;
BEHAVIOUR neighbourSystemType-B BEHAVIOUR
DEFINED AS
"The type of the neighbour system one of: Unknown, End system, Intermediate system, L1 Intermediate system, L2 Intermediate system";;
REGISTERED AS {ISIS.aoi neighbourSystemType (80)};

--11.2.7 Virtual adjacency managed object

virtualAdjacency MANAGED OBJECT CLASS

DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992":top;
CHARACTERIZED BY virtualAdjacency-P PACKAGE
ATTRIBUTES
networkEntityTitle GET,
metric GET;
;;
REGISTERED AS {ISIS.moi virtualAdjacency (2)};

--11.2.7.1 Name bindings for the virtual adjacency managed object

virtualAdjacency-cLNS NAME BINDING

SUBORDINATE OBJECT CLASS virtualAdjacency AND SUBCLASSES;
NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":cLNS AND SUBCLASSES;
WITH ATTRIBUTE networkEntityTitle;
BEHAVIOUR virtualAdjacency-cLNS-B BEHAVIOUR
DEFINED AS This name binding is only applicable where the superior object has an iSType of Level2;;
REGISTERED AS {ISIS.nboi virtualAdjacency-cLNS (3)};

--11.2.7.2 Attributes of the virtual adjacency managed object

metric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.PathMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR metric-B BEHAVIOUR
DEFINED AS Cost of least cost L2 path(s) to destination area based on the default metric;;
REGISTERED AS {ISIS.aoi metric (89)};

--11.2.8 Destination managed object

- This MO class is never instantiated. It exists only to allow the destinationSystem and destinationArea
- MO classes to be derived from it.

destination MANAGED OBJECT CLASS

DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992".top;

CHARACTERIZED BY destination-P PACKAGE

ATTRIBUTES

defaultMetricPathCost GET,
 defaultMetricOutputAdjacencies GET,
 delayMetricPathCost GET,
 delayMetricOutputAdjacencies GET,
 expenseMetricPathCost GET,
 expenseMetricOutputAdjacencies GET,
 errorMetricPathCost GET,
 errorMetricOutputAdjacencies GET;

::

REGISTERED AS {ISIS.moi destination (3)};

--11.2.9 Attributes of the destination managed object**defaultMetricOutputAdjacencies ATTRIBUTE**

WITH ATTRIBUTE SYNTAX ISIS.OutputAdjacencies;

MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;

BEHAVIOUR defaultMetricOutputAdjacencies-B BEHAVIOUR

DEFINED AS

The set of Adjacency (or Reachable Address) managed object identifiers representing the forwarding decisions based upon the default metric for the destination;;

REGISTERED AS {ISIS.aoi defaultMetricOutputAdjacencies (91)};

defaultMetricPathCost ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.PathMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR defaultMetricPathCost-B BEHAVIOUR

DEFINED AS Cost of least cost path(s) using the default metric to destination;;

REGISTERED AS {ISIS.aoi defaultMetricPathCost (90)};

delayMetricOutputAdjacencies ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.OutputAdjacencies;

MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;

BEHAVIOUR delayMetricOutputAdjacencies-B BEHAVIOUR

DEFINED AS

The set of Adjacency (or Reachable Address) managed object identifiers representing the forwarding decisions based upon the delay metric for the destination;;

REGISTERED AS {ISIS.aoi delayMetricOutputAdjacencies (93)};

delayMetricPathCost ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.PathMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR delayMetricPathCost-B BEHAVIOUR

DEFINED AS Cost of least cost path(s) using the delay metric to destination;;

REGISTERED AS {ISIS.aoi delayMetricPathCost (92)};

errorMetricOutputAdjacencies ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.OutputAdjacencies;

MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;

BEHAVIOUR errorMetricOutputAdjacencies-B BEHAVIOUR

DEFINED AS

The set of Adjacency (or Reachable Address) managed object identifiers representing the forwarding decisions based upon the error metric for the destination;;

REGISTERED AS {ISIS.aoi errorMetricOutputAdjacencies (97)};

errorMetricPathCost ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.PathMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR errorMetricPathCost-B BEHAVIOUR

DEFINED AS Cost of least cost path(s) using the error metric to destination;;

REGISTERED AS {ISIS.aoi errorMetricPathCost (96)};

expenseMetricOutputAdjacencies ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.OutputAdjacencies;

MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;

BEHAVIOUR expenseMetricOutputAdjacencies-B BEHAVIOUR

DEFINED AS

The set of Adjacency (or Reachable Address) managed object identifiers representing the forwarding decisions based upon the expense metric for the destination;;

REGISTERED AS {ISIS.aoi expenseMetricOutputAdjacencies (95)};

expenseMetricPathCost ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.PathMetric;

MATCHES FOR EQUALITY, ORDERING;

BEHAVIOUR expenseMetricPathCost-B BEHAVIOUR

DEFINED AS Cost of least cost path(s) using the expense metric to destination;;

REGISTERED AS {ISIS.aoi expenseMetricPathCost (94)};

--11.2.10 Destination system managed object

destinationSystem MANAGED OBJECT CLASS

DERIVED FROM destination;

CHARACTERIZED BY destinationSystem-P PACKAGE

ATTRIBUTES

networkEntityTitle GET;

::

REGISTERED AS {ISIS.moi destinationSystem (4)};

--11.2.10.1 Name bindings

destinationSystem-cLNS NAME BINDING

SUBORDINATE OBJECT CLASS destinationSystem AND SUBCLASSES;

NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":cLNS AND SUBCLASSES;

WITH ATTRIBUTE networkEntityTitle;

REGISTERED AS {ISIS.nboi destinationSystem-cLNS (1)};

--11.2.11 Destination area managed object

destinationArea MANAGED OBJECT CLASS

DERIVED FROM destination;

CHARACTERIZED BY destinationArea-P PACKAGE

ATTRIBUTES

addressPrefix GET; ::

REGISTERED AS {ISIS.moi destinationArea (7)};

--11.2.11.1 Name bindings

destinationArea-cLNS NAME BINDING

SUBORDINATE OBJECT CLASS destinationArea AND SUBCLASSES;
 NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":cLNS AND SUBCLASSES;
 WITH ATTRIBUTE addressPrefix;
 BEHAVIOUR destinationArea-cLNS-B BEHAVIOUR
 DEFINED AS This name binding is only applicable where the superior object has an iSType of Level2;;
 REGISTERED AS {ISIS.nboi destinationArea-cLNS (2)};

--11.2.12 Reachable address managed object

-- Created either through the reachableAddress-linkage-imported
 -- Name binding for reachable addresses instantiated by
 -- the importation of routing information from another
 -- routing protocol (such as the interdomain routing protocol), or the
 -- reachableAddress-linkage-management name binding for adjacencies created
 -- via explicit system management operation.

reachableAddress MANAGED OBJECT CLASS

DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992":top;
 CHARACTERIZED BY reachableAddress-P PACKAGE
 BEHAVIOUR "ISO/IEC 10733":commonCreationDeletion-B, "ISO/IEC 10733":commonStateChange-B;
 ATTRIBUTES
 reachableAddressId GET,
 addressPrefix
 INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B
 GET,
 mappingType
 INITIAL VALUE DERIVATION RULE supplyValueOnCreate-B
 GET,
 defaultMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.defaultMetric-Default
 PERMITTED VALUES ISIS.DefaultMetric-Permitted
 GET-REPLACE,
 delayMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,
 expenseMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,
 errorMetric
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.optionalMetric-Default
 GET-REPLACE,
 defaultMetricType
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.metricType-Default
 GET-REPLACE,
 delayMetricType
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.metricType-Default
 GET-REPLACE,
 expenseMetricType
 REPLACE-WITH-DEFAULT
 DEFAULT VALUE ISIS.metricType-Default

```

    GET-REPLACE,
    errorMetricType
    REPLACE-WITH-DEFAULT
    DEFAULT VALUE ISIS.metricType-Default
    GET-REPLACE,
    "Rec. X.721 | ISO/IEC 10165-2 : 1992":operationalState GET;
ACTIONS
    "Rec. X.723 | ISO/IEC 10165-5":activate,
    "Rec. X.723 | ISO/IEC 10165-5":deactivate;
NOTIFICATIONS
    "Rec. X.721 | ISO/IEC 10165-2 : 1992":objectCreation,
    "Rec. X.721 | ISO/IEC 10165-2 : 1992":objectDeletion,
    "Rec. X.721 | ISO/IEC 10165-2 : 1992":stateChange;
;;
CONDITIONAL PACKAGES
    explicitSNPA-P PRESENT IF the value of mappingType is "explicit",
    extractDSP-P PRESENT IF the value of mappingType is "extractDSP";
REGISTERED AS {ISIS.moi reachableAddress (8)};

```

--11.2.12.1 Explicit SNPA package

```

explicitSNPA-P PACKAGE
    ATTRIBUTES
        sNPAddresses GET-REPLACE;
    REGISTERED AS {ISIS.poi explicitSNPA-P (22)};

```

--11.2.12.2 Extract DSP package

```

extractDSP-P PACKAGE
    BEHAVIOUR extractDSP-P-B BEHAVIOUR
    DEFINED AS
        When present, the remote SNPA address is determined by extracting the bits from the effective NSAP address
        indicated by 1's in the sNPAMask and concatenating them to the sNPAPrefix.;;
    ATTRIBUTES
        sNPAMask
            REPLACE-WITH-DEFAULT
            DEFAULT VALUE ISIS.sNPAMask-Default
            GET-REPLACE,
        sNPAPrefix
            REPLACE-WITH-DEFAULT
            DEFAULT VALUE ISIS.sNPAPrefix-Default
            GET-REPLACE;
    REGISTERED AS {ISIS.poi extractDSP-P (23)};

```

--11.2.12.3 Name bindings

```

reachableAddress-linkage-imported NAME BINDING
    SUBORDINATE OBJECT CLASS reachableAddress AND SUBCLASSES;
    NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":linkage AND SUBCLASSES;
    WITH ATTRIBUTE reachableAddressId;
    BEHAVIOUR reachableAddress-linkage-imported-B BEHAVIOUR
    DEFINED AS
        This name binding is only applicable where the superior object of the linkage instance is an object with iSType
        level2IS It is used for the automatic creation of reachable address MOs. This is useful when injecting intra-domain
        routes obtained through the operation of an interdomain routing protocol.;;
    REGISTERED AS {ISIS.nboi reachableAddress-linkage-imported (5)};

```

```

reachableAddress-linkage-management NAME BINDING
    SUBORDINATE OBJECT CLASS reachableAddress AND SUBCLASSES;
    NAMED BY SUPERIOR OBJECT CLASS "ISO/IEC 10733":linkage AND SUBCLASSES;

```

WITH ATTRIBUTE reachableAddressId;
 BEHAVIOUR reachableAddress-linkage-management-B BEHAVIOUR
 DEFINED AS
 This name binding is only applicable where the superior object of the linkage instance is an object with iSType level2IS It is used for the manual creation of reachable address MOs by system management.;;
 reachableAddressId-B;
 CREATE WITH-REFERENCE-OBJECT reservedName;
 DELETE ONLY-IF-NO-CONTAINED-OBJECTS;
 REGISTERED AS {ISIS.nboi reachableAddress-linkage-management (6)};

--11.2.12.4 Attributes of the reachable address managed object

defaultMetric ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR defaultMetric-B BEHAVIOUR
 DEFINED AS
 The default metric value for reaching the specified prefix over this Circuit. If this attribute is changed while both the Reachable Address and the linkage are Enabled (i.e. state On), the actions described in clause 8.3.5.4 must be taken. ;;
 REGISTERED AS {ISIS.aoi defaultMetric (99)};

defaultMetricType ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.MetricType;
 MATCHES FOR EQUALITY;
 BEHAVIOUR defaultMetricType-B BEHAVIOUR
 DEFINED AS Indicates whether the default metric is internal or external;;
 REGISTERED AS {ISIS.aoi defaultMetricType (103)};

delayMetric ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR delayMetric-B BEHAVIOUR
 DEFINED AS
 The delay metric value for reaching the specified prefix over this Circuit. If this attribute is changed while both the Reachable Address and the linkage are Enabled (i.e. state On), the actions described in clause 8.3.5.4 must be taken.
 The value of zero is reserved to indicate that this metric is not supported;;
 REGISTERED AS {ISIS.aoi delayMetric (100)};

delayMetricType ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.MetricType;
 MATCHES FOR EQUALITY;
 BEHAVIOUR delayMetricType-B BEHAVIOUR
 DEFINED AS Indicates whether the delay metric is internal or external;;
 REGISTERED AS {ISIS.aoi delayMetricType (104)};

errorMetric ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
 MATCHES FOR EQUALITY, ORDERING;
 BEHAVIOUR errorMetric-B BEHAVIOUR
 DEFINED AS
 The error metric value for reaching the specified prefix over this Circuit. If this attribute is changed while both the Reachable Address and the linkage are Enabled (i.e. state On), the actions described in clause 8.3.5.4 must be taken.
 The value of zero is reserved to indicate that this metric is not supported;;
 REGISTERED AS {ISIS.aoi errorMetric (102)};

errorMetricType ATTRIBUTE
 WITH ATTRIBUTE SYNTAX ISIS.MetricType;
 MATCHES FOR EQUALITY;
 BEHAVIOUR errorMetricType-B BEHAVIOUR
 DEFINED AS Indicates whether the error metric is internal or external;;
 REGISTERED AS {ISIS.aoi errorMetricType (106)};

expenseMetric ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.HopMetric;
MATCHES FOR EQUALITY, ORDERING;
BEHAVIOUR expenseMetric-B BEHAVIOUR
DEFINED AS

The expense metric value for reaching the specified prefix over this Circuit. If this attribute is changed while both the Reachable Address and the linkage are Enabled (i.e. state On), the actions described in clause 8.3.5.4 must be taken.

The value of zero is reserved to indicate that this metric is not supported;;

REGISTERED AS {ISIS.aoi expenseMetric (101)};

expenseMetricType ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.MetricType;
MATCHES FOR EQUALITY;
BEHAVIOUR expenseMetricType-B BEHAVIOUR

DEFINED AS Indicates whether the expense metric is internal or external;;

REGISTERED AS {ISIS.aoi expenseMetricType (105)};

mappingType ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.MappingType;
MATCHES FOR EQUALITY;
BEHAVIOUR mappingType-B BEHAVIOUR
DEFINED AS

The type of mapping to be employed to ascertain the SNPA Address which should be used in forwarding NPDUs for this Reachable Address Prefix. The following values of mappingType are defined: none: The mapping is null because the neighbour SNPA is implicit by nature of the subnetwork (e.g. a point-to-point linkage). explicit: The set of subnetwork addresses in the sNPAAAddresses attribute are to be used. extractIDI: The SNPA is embedded in the IDI of the destination NSAP address. The mapping algorithm extracts the SNPA to be used according to the format and encoding rules of ISO8348/Add2. This SNPA extraction algorithm can be used in conjunction with Reachable Address Prefixes from the X.121, F.69, E.163, and E.164 addressing subdomains. extractDSP: All, or a suffix, of the SNPA is embedded in the DSP of the destination address. This SNPA extraction algorithm extracts the embedded subnetwork addressing information by performing a logical AND of the sNPAMask attribute with the destination address. The part of the SNPA extracted from the destination NSAP is appended to the sNPAPrefix to form the next hop subnetwork addressing information;;

REGISTERED AS {ISIS.aoi mappingType (107)};

reachableAddressId ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.GraphicString;
MATCHES FOR EQUALITY, SUBSTRINGS;
BEHAVIOUR reachableAddressId-B BEHAVIOUR
DEFINED AS

A string which is the Identifier for the ReachableAddress and which is unique amongst the set of reachable addresses maintained for this linkage. If this is a reachableAddress created by system management, it is set by the System Manager when the ReachableAddress is created, otherwise it is generated by the implementation such that it is unique. The set of identifiers containing the leading string "Auto" are reserved for reachable addresses imported from other routing protocols. An attempt by system management to create a reachableAddress with such an identifier will cause a reserved name violation;;

REGISTERED AS {ISIS.aoi reachableAddressId (121)};

sNPAMask ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.NAddress;
MATCHES FOR EQUALITY;
BEHAVIOUR sNPAMask-B BEHAVIOUR
DEFINED AS

A Bit mask with 1 bits indicating the positions in the effective destination address from which embedded SNPA information is to be extracted. For the extraction the first octet of the sNPAMask is aligned with the first octet (AFI) of the NSAP Address. If the sNPAMask and NSAP Address are of different lengths, the shorter of the two is logically padded with zeros before performing the extraction;;

REGISTERED AS {ISIS.aoi sNPAMask (122)};

sNPAPrefix ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SNPAPrefix;
 MATCHES FOR EQUALITY;
 BEHAVIOUR sNPAPrefix-B BEHAVIOUR
 DEFINED AS

A fixed SNPA prefix manually configured as an attribute of a Reachable Address with mappingType extractDSP. The SNPA address to use is formed by concatenating the fixed sNPAPrefix with a variable SNPA part that is extracted from the effective destination address. For Reachable Address Prefixes in which the entire SNPA is embedded in the DSP the sNPAPrefix shall be null;;

REGISTERED AS {ISIS.aoi sNPAPrefix (123)};

sNPAAAddresses ATTRIBUTE

WITH ATTRIBUTE SYNTAX ISIS.SNPAAAddresses;
 MATCHES FOR EQUALITY, SET-COMPARISON, SET-INTERSECTION;
 BEHAVIOUR sNPAAAddresses-B BEHAVIOUR
 DEFINED AS

A set of SNPA addresses to which pdu may be forwarded in order to reach an address which matches the address prefix of the Reachable Address.;;

REGISTERED AS {ISIS.aoi sNPAAAddresses (109)};

--11.3 ASN1 modules

ISIS {joint-iso-itu-t network-layer(13) management(0) iSIS(1) asn1Module(2) 0}

DEFINITIONS ::= BEGIN

--- object identifier definitions for identifier prefixes -

isisoi OBJECT IDENTIFIER ::= {NLM.nl iSIS(1)}

sseoi OBJECT IDENTIFIER ::= {isisoi standardSpecificExtensions(0)}

moi OBJECT IDENTIFIER ::= {isisoi managedObjectClass (3)}

poi OBJECT IDENTIFIER ::= {isisoi package (4)}

proi OBJECT IDENTIFIER ::= {isisoi parameter (5)}

nboi OBJECT IDENTIFIER ::= {isisoi nameBinding (6)}

aoi OBJECT IDENTIFIER ::= {isisoi attribute (7)}

agoi OBJECT IDENTIFIER ::= {isisoi attributeGroup (8)}

acoi OBJECT IDENTIFIER ::= {isisoi action (9)}

noi OBJECT IDENTIFIER ::= {isisoi notification (10)}

---object identifiers for notification parameters -

se OBJECT IDENTIFIER ::= {sseoi specificProblems(3)}

areaMismatch OBJECT IDENTIFIER ::= {se areaMismatch(0)}

attemptToExceedMaximumSequenceNumber OBJECT IDENTIFIER ::=

{se attemptToExceedMaximumSequenceNumber(1)}

authenticationFailure OBJECT IDENTIFIER ::= {se authenticationFailure(2)}

corruptedLSPsDetected OBJECT IDENTIFIER ::= {se corruptedLSPsDetected(3)}

iDFieldLengthMismatch OBJECT IDENTIFIER ::= {se iDFieldLengthMismatch(4)}

lanL1DesignatedIntermediateSystemChange OBJECT IDENTIFIER ::=

{se lanL1DesignatedIntermediateSystemChange(5)}

ISPL1DatabaseOverload OBJECT IDENTIFIER ::= {se ISPL1DatabaseOverload(6)}

ISPL2DatabaseOverload OBJECT IDENTIFIER ::= {se ISPL2DatabaseOverload(7)}

manualAddressDroppedFromArea OBJECT IDENTIFIER ::= {se manualAddressDroppedFromArea(8)}

maximumAreaAddressesMismatch OBJECT IDENTIFIER ::= {se maximumAreaAddressesMismatch(9)}

ownLSPPurge OBJECT IDENTIFIER ::= {se ownLSPPurge(10)}

partitionVirtualLinkChange OBJECT IDENTIFIER ::= {se partitionVirtualLinkChange(11)}

rejectedAdjacency OBJECT IDENTIFIER ::= {se rejectedAdjacency(12)}

sequenceNumberSkip OBJECT IDENTIFIER ::= {se sequenceNumberSkip(13)}

versionSkew OBJECT IDENTIFIER ::= {se versionSkew(14)}

--11.3.1 ASN.1 types and values

AddressPrefix ::= BIT STRING(SIZE(0..160)) -- Size shall be a multiple of four, since the protocol represents these as semi-octets

AdjacencyState ::= ENUMERATED { initializing(0), up(1), failed(2), down(3) }

AdjacencyUsage ::= ENUMERATED{ undefined(0), level1(1), level2(2), level1and2(3)}
 AreaAddress ::= OCTET STRING(SIZE(1..20))
 AreaAddresses ::= SET OF AreaAddress
 Boolean ::= BOOLEAN
 CircuitID ::= OCTET STRING(SIZE(2..9))
 CircuitType ::= ENUMERATED{ broadcast(0), ptToPt(1), staticIn(2), staticOut(3), dA(4)}
 DatabaseState ::= ENUMERATED{ off(0), on(1), waiting(2)}
 DesignatedISChange ::= ENUMERATED{ resigned(0), elected(1)}
 HopMetric ::= INTEGER(0..maxLinkMetric)
 IDLength ::= INTEGER(0..8)
 IntermediateSystemPriority ::= INTEGER(1..127)
 ISType ::= ENUMERATED{ level1IS(1), level2IS(2)}
 LocalDistinguishedName ::= CMIP-1.BaseManagedObjectId
 LSPID ::= OCTET STRING(SIZE(2..11))
 MappingType ::= ENUMERATED{ none(0), explicit(1), extractIDI(2), extractDSP(3)}
 MaximumAreaAddresses ::= INTEGER(0..254)
 MaximumPathSplits ::= INTEGER(1..32)
 MaximumVirtualAdjacencies ::= INTEGER(0..32)
 maxLinkMetric INTEGER ::= 63
 maxPathMetric INTEGER ::= 1023
 MetricType ::= ENUMERATED{ internal(0), external(1)}
 NAddress ::= NLM.NAddress
 NeighbourSystemType ::= ENUMERATED{ unknown(0), endSystem(1), intermediateSystem(2), l1IntermediateSystem(3), l2IntermediateSystem(4)}
 Null ::= NULL
 ObjectIdentifier ::= OBJECT IDENTIFIER
 OctetString ::= OCTET STRING
 OriginatingLSPBufferSize ::= INTEGER(512..1492)
 originatingLSPBufferSizeMismatch OBJECT IDENTIFIER ::= {se originatingLSPBufferSizeMismatch(15)}
 OutputAdjacencies ::= SET OF LocalDistinguishedName
 OverloadStateChange ::= ENUMERATED{ on(0), waiting(1)}
 Password ::= OCTET STRING(SIZE(0..254))
 Passwords ::= SET OF Password
 PathMetric ::= INTEGER(0..maxPathMetric)
 Reason ::= ENUMERATED{ holdingTimerExpired(0), checksumError(1), oneWayConnectivity(2), callRejected(3), reserveTimerExpired(4), circuitDisabled(5), versionSkew(6), areaMismatch(7), maximumBroadcastIntermediateSystemsExceeded(8), maximumBroadcastEndSystemsExceeded(9), wrongSystemType(10)}
 ReceiveLSPBufferSize INTEGER ::= 1492
 SNPAAddress ::= NLM.SNPAAddress
 SNPAAddresses ::= SET OF SNPAAddress
 SNPAPrefix ::= BIT STRING(SIZE(0..120))
 sNPAPrefix-Default SNPAPrefix ::= 'B
 sNPAMask-Default NAddress ::= 'B
 SourceId ::= OCTET STRING(SIZE(1..10))
 SystemId ::= OCTET STRING(SIZE(0..8))
 SystemIds ::= SET OF SystemId
 Version ::= GraphicString
 VirtualLinkChange ::= ENUMERATED{ deleted(0), created(1)}
 zero INTEGER ::= 0

--11.3.2 Defaults and permitted values

callEstablishmentMetricIncrement-Default INTEGER ::= 0
 completeSNPInterval-Default CommonMgt.Timer ::= {exponent 0, mantissa 10}
 defaultMetric-Default INTEGER ::= 20
 DefaultMetric-Permitted ::= INTEGER(1..maxLinkMetric)
 dRISISHelloTimer-Default CommonMgt.Timer ::= {exponent 0, mantissa 1}
 externalDomain-Default BOOLEAN ::= TRUE
 iSISHelloTimer-Default CommonMgt.Timer ::= {exponent 0, mantissa 3}
 l1IntermediateSystemPriority-Default INTEGER ::= 64

l2IntermediateSystemPriority-Default INTEGER ::= 64
 manualAreaAddresses-Default AreaAddresses ::= {}
 manualL2OnlyMode-Default BOOLEAN ::= FALSE
 maximumAreaAddresses-Default INTEGER ::= 3
 maximumPathSplits-Default INTEGER ::= 2
 maximumLSPGenerationInterval-Default CommonMgt.Timer ::= {exponent 2, mantissa 9}
 maximumVirtualAdjacencies-Default INTEGER ::= 2
 metricType-Default MetricType ::= internal
 minimumBroadcastLSPTransmissionInterval-Default CommonMgt.Timer ::= {exponent -3, mantissa 33}
 minimumLSPGenerationInterval-Default CommonMgt.Timer ::= {exponent 0, mantissa 30}
 minimumLSPTransmissionInterval-Default CommonMgt.Timer ::= {exponent 0, mantissa 5}
 neighbourSNPAAddress-Default SNPAAddress ::= {type NLM.sNPADTEAddress, address 'H}
 optionalMetric-Default INTEGER ::= 0
 originatingL1LSPBufferSize-Default INTEGER ::= receiveLSPBufferSize
 originatingL2LSPBufferSize-Default INTEGER ::= receiveLSPBufferSize
 partialSNPInterval-Default CommonMgt.Timer ::= {exponent 0, mantissa 2}
 password-Default Password ::= 'H
 passwords-Default Passwords ::= {} -- The empty set
 pollESHelloRate-Default CommonMgt.Timer ::= {exponent 0, mantissa 50}
 reserveTimer-Default CommonMgt.Timer ::= {exponent 2, mantissa 6}
 sNPAAddresses-Default SNPAAddresses ::= {}
 waitingTime-Default CommonMgt.Timer ::= {exponent 0, mantissa 60}

END

12 Conformance

12.1 Conformance for protocol implementation

This subclause specifies the conformance for protocol implementation of intra-domain routing protocol. The supplier of implementation for protocol implementation shall support the following specification.

NOTE 72 The conformance for protocol implementation is independent of the conformance for management information implementation specified in subclause 12.2.

12.1.1 Static conformance requirements

12.1.1.1 Protocol implementation conformance statement

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

12.1.1.2 Static conformance requirements for all ISs

A system claiming conformance to this International Standard shall be capable of:

- a) calculating a single minimum cost route to each destination according to 7.2.6 for the default metric specified in 7.2.2;
- b) utilising Link State information from a system only when an LSP with LSP number 0 and remaining lifetime>0 is present according to 7.2.5;
- c) selection of paths according to 7.2.7 and 7.2.12
- d) performing the robustness checks according to 7.2.8;
- e) constructing a forwarding database according to 7.2.9;

- f) if (and only if) Area Partition Repair is supported,
 - 1) performing the operations according to 7.2.10;
 - 2) performing the encapsulation operations in the forwarding process according to 7.4.3.2; and
 - 3) performing the decapsulation operations in the receive process according to 7.4.4.2;
- g) assigning and computing area addresses according to 7.1.5 and 7.2.11;
- h) generating local link state information as required by 7.3.2;
- i) including information from Manual Adjacencies according to 7.3.3.1;
- j) if (and only if) reachable addresses are supported, including information from reachable addresses according to 7.3.3.2;
- k) generating multiple LSPs according to 7.3.4;
- l) generating LSPs periodically according to 7.3.5;
- m) generating LSPs on the occurrence of events according to 7.3.6;
- n) generating an LSP checksum according to 7.3.11;
- o) operating the Update process according to 7.3.12–7.3.17 including controlling the rate of LSP transmission only for each broadcast circuit (if any) according to 7.3.15.6;
- p) operating the LSP database overload procedures according to 7.3.19.1;
- q) selecting the appropriate forwarding database according to 7.4.2;
- r) forwarding ISO 8473 PDUs according to 7.4.3.1 and 7.4.3.3;
- s) operating the receive process according to 7.4.4.1;
- t) performing on each supported point-to-point circuit (if any):
 - 1) forming and maintaining adjacencies according to 8.2;
- u) performing on each supported ISO 8208 circuit (if any)
 - 1) SVC establishment according to 8.3.2.1 using the network layer protocols according to 8.3.1;
 - 2) If reachable addresses are supported, the operations specified in 8.3.2.2 – 8.3.5.6.
 - 3) If **callEstablishmentMetricIncrement** greater than zero are supported, the operations specified in 8.3.5.3.
 - 4) If the reverse path cache is supported, the operations specified in 8.3.3
- v) performing on each supported broadcast circuit (if any)
 - 1) the pseudonode operations according to 7.2.3;
 - 2) controlling the rate of LSP transmission according to 7.3.15.6;
 - 3) the operations specified in 8.4.2–8.4.5 and 8.4.7;
 - 4) the operations specified in 8.4.5.
- w) constructing and correctly parsing all PDUs according to clause 9;
- x) providing a system environment in accordance with clause 10;
- y) being managed via the system management attributes defined in clause 11. For all attributes referenced in the normative text, the default value (if any) shall be supported. Other values shall be supported if referenced in a REQUIRED VALUES clause of the GDMO definition;
- z) If authentication procedures are implemented:
 - 1) the authentication field processing functions of clauses 7.3.7–7.3.10, 7.3.15.1–7.3.15.4, 8.2.3–8.2.4, and 8.4.2.1;
 - 2) the **Authentication Information** field of the PDU in clauses 9.5–9.13.

12.1.1.3 Static conformance requirements for level 1 ISs

A system claiming conformance to this International Standard as a level 1 IS shall conform to the requirements of 12.1.2 and in addition shall be capable of

- a) identifying the nearest Level 2 IS according to 7.2.9.1;
- b) generating Level 1 LSPs according to 7.3.7;
- c) generating Level 1 pseudonode LSPs for each supported broadcast circuit (if any) according to 7.3.8;
- d) performing the actions in Level 1 Waiting State according to 7.3.19.2

12.1.1.4 Static conformance requirements for level 2 ISs

A system claiming conformance to this International Standard as a level 2 IS shall conform to the requirements of 12.1.2 and in addition shall be capable of

- a) setting the attached flag according to 7.2.9.2;
- b) generating Level 2 LSPs according to 7.3.9;
- c) generating Level 2 pseudonode LSPs for each supported broadcast circuit (if any) according to 7.3.10;
- d) performing the actions in Level 2 Waiting State according to 7.3.19.3.

12.1.2 Dynamic conformance**12.1.2.1 Receive process conformance requirements**

Any protocol function supported shall be implemented in accordance with 7.4.4.

12.1.2.2 Update process conformance requirements

Any protocol function supported shall be implemented in accordance with 7.3 and its subclauses.

Any PDU transmitted shall be constructed in accordance with the appropriate subclauses of clause 9.

12.1.2.2.1 Decision process conformance requirements

Any protocol function supported shall be implemented in accordance with 7.2 and its subclauses.

12.1.2.3 Forwarding process conformance requirements

Any protocol function supported shall be implemented in accordance with 7.4 and its subclauses.

12.1.2.4 Performance requirements

This International Standard requires that the following performance criteria be met. These requirements apply regardless of other demands on the system; if an Intermediate system has other tasks as well, those will only get resources not required to meet these criteria.

Each Intermediate system implementation shall specify (in its PICS):

- a) the maximum number of other Intermediate systems it can handle. (For L1 Intermediate systems that means Intermediate systems in the area; for L2 Intermediate systems that is the sum of Intermediate systems in the area and Intermediate systems in the L2 subdomain.) Call this limit N.
- b) the maximum supported forwarding rate in ISO 8473 PDUs per second.

12.1.2.4.1 Performance requirements on the update process

The implementation shall guarantee the update process enough resources to process N LSPs per 30 seconds. (Resources = CPU, memory, buffers, etc.)

In a stable topology the arrival of a single new LSP on a circuit shall result in the propagation of that new LSP over the other circuits of the IS within one second, irrespective of the forwarding load for ISO 8473 data PDUs.

12.1.2.4.2 Performance requirements on the decision process

The implementation shall guarantee the decision process enough resources to complete (i.e. start to finish) within 5 seconds, in a stable topology while forwarding at the maximum rate. (For L2 Intermediate Systems, this applies to the two levels together, not each level separately.)

12.1.2.4.3 Reception and processing of PDUs

An ideal Intermediate system would be able to correctly process all PDUs, both control and data, with which it was presented, while simultaneously running the decision process and responding to management requests. However, in the implementations of real Intermediate systems some compromises must be made. The way in which these compromises are made can dramatically affect the correctness of operation of the Intermediate system. The following general principles apply.

- a) A stable topology should result in stable routes when forwarding at the maximum rated forwarding rate.
- b) Some forwarding progress should always be made (albeit over incorrect routes) even in the presence of a maximally unstable topology.

In order to further characterise the required behaviour, it is necessary to identify the following types of traffic.

- c) IIH traffic. This traffic is important for maintaining Intermediate system adjacencies and hence the Intermediate system topology. In order to prevent gratuitous topology changes it is essential that Intermediate system adjacencies are not caused to go down erroneously. In order to achieve this no more than $\text{ISISHoldingMultiplier} - 1$ IIH PDUs may be dropped between any pair of Intermediate systems. A safer requirement is that *no* IIH PDUs are dropped.

The rate of arrival of IIH PDUs is approximately constant and is limited on point-to-point links to $1/\text{ISISHelloTimer}$ and on LANs to a value of approximately $2(n/\text{ISISHelloTimer}) + 2$, where n is the number of Intermediate systems on the LAN (assuming the worst case that they are all Level 2 Intermediate systems).

- d) ESH PDU traffic. This traffic is important for maintaining End system adjacencies, and has relatively low processing latency. As with IIH PDUs, loss of End system adjacencies will cause gratuitous topology changes which will result in extra control traffic.

The rate of arrival of ESH PDUs on point-to-point links is limited to approximately $1/\text{DefaultESHelloTimer}$ under all conditions. On LANs the background rate is approximately $n/\text{DefaultESHelloTimer}$ where n is the number of End systems on the LAN. The maximum rate during polling is limited to approximately $n/\text{pollIES-HelloRate}$ averaged over a period of about 2 minutes. (Note that the actual peak arrival rate over a small interval may be much higher than this.)

- e) LSP (and SNP) traffic. This traffic will be retransmitted indefinitely by the update process if it is dropped, so there is no requirement to be able to process every received PDU. However, if a substantial proportion are lost, the rate of convergence to correct routes will be affected, and bandwidth and processing power will be wasted.

On point-to-point links the peak rate of arrival is limited only by the speed of the data link and the other traffic flowing on that link. The maximum average rate is determined by the topology.

On LANs the rate is limited at a first approximation to a maximum rate of $1000/\text{minimumBroadcastLSP-TransmissionInterval}$, however it is possible that this may be multiplied by a factor of up to n , where n is the number of Intermediate systems on the LAN, for short periods. An Intermediate system shall be able to receive and process at least the former rate without loss, even if presented with LSPs at the higher rate. (i.e. it is permitted to drop LSPs, but must process at least $1000/\text{minimumBroadcastLSPTransmissionInterval}$ per second of those presented.)

The maximum background rate of LSP traffic (for a stable topology) is dependent on the maximum supported configuration size and the settings of $\text{maximumLSP-GenerationInterval}$. For these purposes the default value of 900 seconds can be assumed. The number of LSPs per second is then very approximately $(n_1 + n_2 + n_e/x)/900$ where n_1 is the

number of level 1 Intermediate systems, n_2 the number of level 2 Intermediate systems, n_e the number of End system IDs and x the number of ID which can be fitted into a single LSP.

NOTE 73 This gives a value around 1 per second for typical maximum configurations of:

4000 IDs
100 L1 Intermediate systems per area
400 L2 Intermediate systems.

- f) Data Traffic. This is theoretically unlimited and can arrive at the maximum data rate of the point-to-point link or LAN (e.g. 14 000 PDUs per second for a 10 Mbps CSMA/CD LAN). In practice it will be limited by the operation of the congestion avoidance and control algorithms, but owing to the relatively slow response time of these algorithms, substantial peaks are likely to occur.

An Intermediate system shall state in its PICS its maximum forwarding rate. This shall be quoted under at least the following conditions.

- 1) A stable topology of maximum size.
- 2) A maximally unstable topology. This figure shall be non-zero, but may reasonably be as low as 1 PDU per second.

The following constraints must be met.

- g) The implementation shall be capable of receiving the maximum rate of ISH PDUs without loss whenever the following conditions hold:
- 1) the data forwarding traffic rate averaged over any period of one second does not exceed the rate which the implementation claims to support; and
 - 2) the ESH and LSP rates do not exceed the background (stable topology) rate.
- h) If it is unavoidable that PDUs are dropped, it is a goal that the order of retaining PDUs shall be as follows (i.e. It is least desirable for IIH PDUs to be dropped).
- 1) IIH PDUs
 - 2) ESH PDUs
 - 3) LSPs and SNPs
 - 4) data PDUs.

However, no class of traffic shall be completely starved. One way to achieve this is to allocate a queue of suitable length to each class of traffic and place the PDUs onto the appropriate queue as they arrive. If the queue is full the PDUs are discarded. Processor resources shall be allocated to the queues to ensure that they all make progress with the same priorities as above. This model assumes that an implementation is capable of receiving PDUs and selecting their correct queue at the maximum possible data rate (14 000 PDUs per second for a LAN). If this is not the case, reception of data traffic at a rate greater than some limit (which must be greater than the maximum rated limit) will cause loss of some IIH PDUs even in a stable topology. This limit shall be quoted in the PICS if it exists.

NOTE 74 Starting from the stable topology condition at maximum data forwarding rate, an increase in the arrival rate of data PDUs will initially only cause some data NPDUs to be lost. As the rate of arrival of data NPDUs is further increased a point may be reached at which random PDUs are dropped. This is the rate which must be quoted in the PICS.

12.1.2.4.4 Transmission

Sufficient processor resources shall be allocated to the transmission process to enable it to keep pace with reception for each PDU type. Where prioritisation is required, the same order as for reception of PDU types applies.

12.2 Conformance for management information implementation

This subclause specifies the conformance for management information implementation of intra-domain routing protocol. The supplier of implementation for management information implementation shall support the following specification.

NOTE 75 The conformance for management information implementation is independent of the conformance for protocol implementation specified in subclause 12.1

12.2.1 Static Conformance

The implementation shall conform to the requirements of this International Standard in the manager role, the agent role, or both roles. A claim of conformance to at least one role shall be made in Table G.1 of this International Standard.

If a claim of conformance is made for support in the manager role, the implementation shall support at least one management operation or notification or action of the managed objects specified by this International Standard. The conformance requirements in the manager role for those management operations and notifications and actions are identified in Table G.3 and further tables referenced by Annex G.

If a claim of conformance is made for support in the agent role, the implementation shall support one or more instances of the adjacency managed object class, the destination managed object class, the destination area managed object, the destination system managed object class, the reachable address managed object class and the virtual adjacency managed object class identified in Table G.4 of this International Standard and further tables referenced by Annex G.

If a claim of conformance is made for support in the agent role, the implementation shall support at least one name binding identified in Table G.8 of this International Standard for each supported managed object.

The implementation shall support the transfer syntax derived from the encoding rules specified in ITU-T Rec. X.690 | ISO/IEC 8825-1 named {joint-iso-itu-t asn1(1) basic-encoding(1)} for the abstract data types referenced by the definitions for which support is claimed.

12.2.2 Dynamic conformance

Implementations claiming to conform to this International Standard shall support the elements of procedure and definitions of semantics corresponding to the definitions for which support is claimed.

12.2.3 Management implementation conformance statement requirements

Any MCS proforma, MICS proforma, MOCS proforma, and MRCS proforma which conform to this International Standard shall be technically identical to the proformas specified in Annexes G, G. I. and J preserving table numbering and the index numbers of items, and differing only in pagination and page headers.

The supplier of an implementation which is claimed to conform to this International Standard shall complete a copy of the management conformance summary (MCS) provided in Annex G as part of the conformance requirements together with any other ICS proformas referenced as applicable from that MCS. Any MCS, MICS, MOCS, and MRCS which conform to this International Standard shall:

- describe an implementation which conforms to this International Standard;
- have been completed in accordance with the instructions for completion given in ITU-T Rec. X.724 ISO/IEC 10165-6;
- include the information necessary to uniquely identify both the supplier and the implementation.

12.2.4 IS level specific conformance requirements

The supplier of an implementation which is claimed to conform to this International Standard shall support at least level 1 IS identified in Table G.2 of this International Standard.

12.2.4.1 Conformance to level 1 ISs

An implementation claiming conformance to the level 1 IS of this International Standard in the agent role as a managed implementation shall support the adjacency MO and the destination MO

12.2.4.2 Conformance to level 2 ISs

An implementation claiming conformance to the level 2 IS of this International Standard in the agent role as a managed implementation shall support the adjacency MO, the destination system MO, the destination area MO, the virtual adjacency MO and the reachable address MO.

Annex A (normative)

PICS pro forma

A.1 Introduction

The supplier of a protocol implementation which is claimed to conform to International Standard ISO 10589, whether as a level 1 or level 2 Intermediate system implementation, shall complete the applicable Protocol Implementation Conformance Statement (PICS) pro forma.

A completed PICS pro forma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. The PICS can have a number of uses, including use:

- by the protocol implementor, as a check-list to reduce the risk of failure to conform to the International Standard through oversight;
- by the supplier and acquirer — or potential acquirer — of the implementation, as a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the International Standard PICS pro forma;
- by the user — or potential user — of the implementation, as a basis for initially checking the possibility of interworking with another implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible PICS's);
- by a protocol tester, as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

A.2 Abbreviations and special symbols

A.2.1 Status-related symbols

- M mandatory
 O optional
 O.<n> optional, but support of at least one of the group of options labelled by the same numeral <n> is required.
 X prohibited
 — not applicable
 c.<p> conditional requirement, according to condition <p>

A.3 Instructions for completing the pics pro formas

A.3.1 General structure of the PICS pro forma

The first part of the PICS pro forma — Implementation Identification and Protocol Summary — is to be completed as indicated with the information necessary to identify fully both the supplier and the implementation.

The main part of the PICS pro forma is a fixed-format questionnaire divided into subclauses each containing a group of individual items. Answers to the questionnaire items are to be provided in the rightmost column, either by simply marking an answer to indicate a restricted choice (usually Yes or No), or by entering a value or a set or range of values. (Note that there are some items where two or more choices from a set of possible answers can apply: all relevant choices are to be marked.)

Each item is identified by an item reference in the first column; the second column contains the question to be answered; the third column contains the reference or references to the material that specifies the item in the main body of the International Standard. The remaining columns record the status of the item — whether support is mandatory, optional or conditional — and provide the space for the answers: see A.3.4 below.

A supplier may also provide — or be required to provide — further information, categorised as either Additional Information or Exception Information. When present, each kind of further information is to be provided in a further subclause of items labelled A<i> or X<i> respectively for cross-referencing purposes, where <i> is any unambiguous identification for the item (e.g. simply a number): there are no other restrictions on its format and presentation.

A completed PICS pro forma, including any Additional Information and Exception Information, is the Protocol Implementation Conformance Statement for the implementation in question.

NOTE 76 Where an implementation is capable of being configured in more than one way, a single PICS may be able to describe all such configurations. However, the supplier has the choice of providing more than one PICS, each covering some subset of the implementation's configuration capabilities, in case this makes for easier and clearer presentation of the information.

A.3.2 Additional information

Items of Additional Information allow a supplier to provide further information intended to assist the interpretation of the PICS. It is not intended or expected that a large quantity will be supplied, and a PICS can be considered complete without any such information. Examples might be an outline of the ways in which a (single) implementation can be set up to operate in a variety of environments and configurations.

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

A.3.3 Exception information

It may occasionally happen that a supplier will wish to answer an item with mandatory or prohibited status (after any conditions have been applied) in a way that conflicts with the indicated requirement. No pre-printed answer will be found in the Support column for this, but the Supplier may write the desired answer into the Support column. If this is done, the supplier is required to provide an item of Exception Information containing the appropriate rationale, and a cross-reference from the inserted answer to the Exception item.

An implementation for which an Exception item is required in this way does not conform to ISO 10589.

NOTE 77 A possible reason for the situation described above is that a defect report is being progressed, which is expected to change the requirement that is not met by the implementation.

A.3.4 Conditional status

A.3.4.1 Conditional items

The PICS pro forma contains a number of conditional items. These are items for which the status — mandatory, optional or prohibited — that applies is dependent upon whether or not certain other items are supported, or upon the values supported for other items. In many cases, whether or not the item applies at all is conditional in this way, as well as the status when the item does apply.

Individual conditional items are indicated by a conditional symbol in the Status column as described in A.3.4.2 below. Where a group of items are subject to the same condition for applicability, a separate preliminary question about the condition appears at the head of the group, with an instruction to skip to a later point in the questionnaire if the "Not Applicable" answer is selected.

A.3.4.2 Conditional symbols and conditions

A conditional symbol is of the form $c.<n>$ or $c.G<n>$ where $<n>$ is a numeral. For the first form, the numeral identifies a condition appearing in a list at the end of the subclause containing the item. For the second form, $c.G<n>$, the numeral identifies a condition appearing in the list of global conditions at the end of the PICS.

A simple condition is of the form:

if $<p>$ then $<s1>$ else $<s2>$

where $<p>$ is a predicate (see A.3.4.3 below), and $<s1>$ and

$<s2>$ are either basic status symbols (M, O, X, or —) or the symbol "—". An extended condition is of the form

if $<p1>$ then $<s1>$ else $<s2>$
 else if $<p2>$ then $<s2>$
 [else if $<p3>$...]
 else $<s n >$

where $<p1>$ etc. are predicates and $<s1>$ etc. are basic status symbols or "—".

The status symbol applicable to an item governed by a simple condition is <s1> if the predicate of the condition is true, and <s2> otherwise; the status symbol applicable to an item governed by an extended condition is <s i > where <p i > is the first true predicate, if any, in the sequence <p1>, <p2>..., and <s n > if no predicate is true.

A.3.4.3 Predicates

A simple predicate in a condition is either

- a) a single item reference; or
- b) a relation containing a comparison operator (=, <, etc.) with one (or both) of its operands being an item reference for an item taking numerical values as its answer.

In case (a) the predicate is true if the item referred to is marked as supported, and false otherwise. In case (b), the predicate is true if the relation holds when each item reference is replaced by the value entered in the Support column as answer to the item referred to.

Compound predicates are Boolean expressions constructed by combining simple predicates using the Boolean operators AND, OR and NOT, and parentheses, in the usual way. A compound predicate is true if and only if the Boolean expression evaluates to true when the simple predicates are interpreted as described above.

Items whose references are used in predicates are indicated by an asterisk in the Item column.

A.3.4.4 Answering conditional items

To answer a conditional item, the predicate(s) of the condition is (are) evaluated as described in A.3.4.3 above, and the applicable status symbol is determined as described in A.3.4.2. If the status symbol is “–” this indicates that the item is to be marked in this case; otherwise, the Support column is to be completed in the usual way.

When two or more basic status symbols appear in a condition for an item, the Support column for the item contains one line for each such symbol, labelled by the relevant symbol. The answer for the item is to be marked in the line labelled by the symbol selected according to the value of the condition (unselected lines may be crossed out for added clarity).

For example, in the item illustrated below, the N/A column would be marked if neither predicate were true; the answer line labelled “M:” would be marked if item A4 was marked as supported, and the answer line labelled “O:” would be marked if the condition including items D1 and B52 applied.

Item		References	Status	N/A	Support
H3	Is ... supported?	42.3(d)	C.1	☐	M: Yes ☐ O: Yes ☐ No ☐

C.1 if A4 then M
else if D1 AND (B52 < 3) then O else –

A.4 Identification

A.4.1 Implementation identification

A.4.2 Protocol summary: ISO 10589

Supplier	
Contact point for queries about this PICS	
Implementation Name(s) and Version(s)	
Operating system Name(s) and Version(s)	
Other Hardware and Operating Systems Claimed	
System Name(s) (if different)	

Notes:

- Only the first three items are required for all implementations; others may be completed as appropriate in meeting the requirements for full identification.
- The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (using, e.g., Type, Series, Model)

A.4.3 Protocol summary: ISO 10589

Protocol Version	
Addenda Implemented (if applicable)	
Amendments Implemented	
Have any exception items been required (see A.3.3)? No ☐ Yes ☐ (The answer Yes means that the implementation does not conform to ISO 10589)	
Date of Statement	

A.5 Protocol summary: ISO 10589 general

Item	Functionality/Description	References	Status	Support
AllIS	Are all basic IS-IS routing functions implemented?	12.1.2	M	Yes ☐
System Management	Is the system capable of being managed by the specified management information ?	11	M	Yes ☐
Authentication	Is PDU authentication based on passwords implemented?	7.3.7-7.3.10, 7.3.15.1 – 7.3.15.4, 8.2.3-8.2.4, 8.4.1.1	O	Yes ☐ No ☐
Default Metric	Is the default metric supported?	7.2.2, 7.2.6	M	Yes ☐
Delay Metric	Is the delay metric supported?	7.2.2, 7.2.6	O	Yes ☐ No ☐
Expense Metric	Is the expense metric supported?	7.2.2, 7.2.6	O	Yes ☐ No ☐
Error Metric	Is the error metric supported?	7.2.2, 7.2.6	O	Yes ☐ No ☐
ID Field Length	What values of RouteingDomainIDLength (from the set 1-8) are supported by this implementation? Is the value configurable by system management?	7.1.3	M	values = Yes ☐ No ☐
Forwarding Rate	How many ISO 8473 PDUs can the implementation forward per second ?	12.2.4.1.b	M	PDUs/sec =
Performance	Are the implementation performance criteria met?	12.2.4	M	Yes ☐

A.5.1 System environment: general

Item	Functionality/Description	References	Status	Support
ISO9542	Are the appropriate ISO 9542 operations implemented?	10.3, 8.2.1-8.2.2, 8.3.4, 8.4.5, 8.4.6	M	Yes ☐
Timer Jitter	Is jitter introduced in all periodic timers whose expiration causes transmission of a PDU?	10.1	M	Yes ☐

A.5.2 Subnetwork dependent functions: general

Item	Functionality/Description	References	Status	Support
*LAN	Are the subnetwork dependent functions for broadcast subnetworks implemented?	8.4	O.1	Yes ☐ No ☐
LAN IS Adjacencies	Are the LAN IS adjacency establishment operations implemented?	8.4.1 – 8.4.3	LAN: M	N/A ☐ Yes ☐
LAN ES Adjacencies	Are the LAN ES adjacency establishment operations implemented?	8.4.6	LAN: M	N/A ☐ Yes ☐
LAN DIS	Are the LAN designated IS operations implemented?	8.4.4, 8.4.5	LAN: M	N/A ☐ Yes ☐
*8208 Static	Are the subnetwork dependent functions for ISO 8208 subnetworks implemented?	8.3	O.1	Yes ☐ No ☐
8208 SND CF	Are the ISO8208 Subnetwork Dependent Convergence Functions implemented?	8.3.1, 8.3.2.1	C.1: M	N/A ☐ Yes ☐
*PtPt	Are the subnetwork dependent functions for point-to-point subnetworks implemented?	8.2	O.1	Yes ☐ No ☐
PtPt IS Adjacencies	Are the point-to-point IS adjacency establishment operations implemented?	8.2.2 – 8.2.5	C.2: M	N/A ☐ Yes ☐
PtPt ES Adjacencies	Are the point-to-point ES adjacency establishment operations implemented?	8.2.1	C.2: M	N/A ☐ Yes ☐
PtPt IIH PDU	Are point-to-point IIH PDUs correctly constructed and parsed?	9.7	C.2: M	N/A ☐ Yes ☐

C.1 if 8208 Static or 8208 DA then M else –

C.2 if PtPt or 8208 Static then M else –

A.5.3 Update process: general

Item	Functionality/Description	References	Status	Support
LSP Periodic Generation	Is periodic generation of new local LSPs implemented?	7.3.2, 7.3.5, 7.3.13	M	Yes ☐
LSP Event Driven Generation	Is event driven generation of new local LSPs implemented?	7.3.6	M	Yes ☐
Pseudonode LSP Generation	Is generation of pseudonode LSPs implemented?	7.3.8, 7.3.10	LAN:M	N/A ☐ Yes ☐
Multiple LSP Generation	Is multiple LSP generation implemented?	7.3.4	M	Yes ☐
LSP Propagation	Is propagation of LSPs implemented?	7.3.12, 7.3.14, 7.3.15.1, 7.3.15.5	M	Yes ☐
LSP Lifetime Control	Are the LSP lifetime control operations implemented?	7.3.16.3, 7.3.15.4	M	Yes ☐
CSNP Generation	Is the generation of CSNPs implemented?	7.3.15.3, 7.3.17	M	Yes ☐
PSNP Generation	Is the generation of PSNPs implemented?	7.3.15.4, 7.3.17	M	Yes ☐
SNP Processing	Are the sequence number PDU processing procedures implemented?	7.3.15.2, 7.3.17	M	Yes ☐
LSDB Overload	Are the LSP database overload operations implemented?	7.3.19	M	Yes ☐

A.5.4 Decision process: general

Item	Functionality/Description	References	Status	Support
Minimum Cost Path	Is computation of a single minimum cost path based upon each supported metric implemented?	7.2.6	M	Yes ☐
Equal Cost Paths	Is computation of equal minimum cost paths based upon each supported metric implemented?	7.2.6	O	Yes ☐ No ☐
Down Stream Paths	Is computation of downstream routes based upon each supported metric implemented?	7.2.6	O	Yes ☐ No ☐
Multiple LSPs Recognition	Are multiple LSPs used only when a LSP with LSP#0 and remaining lifetime greater 0 is present?	7.2.5	M	Yes ☐
Overloaded IS Exclusion	Are links to Iss with overloaded LSDBs ignored?	7.2.8.1	M	Yes ☐
Two Way Connectivity	Are links not reported by both end ISs ignored?	7.2.8.2	M	Yes ☐
Path Preference	Is the order of preference for path selection implemented?	7.2.12	M	Yes ☐
Excess Path Removal	Is removal of excess paths implemented?	7.2.7	M	Yes ☐
FIB Construction	Is the construction of ISO8473 Forwarding Information Bases implemented?	7.2.9	M	Yes ☐

A.5.5 Forward/receive process: general

Item	Functionality/Description	References	Status	Support
FIB Selection	Is selection of appropriate Forwarding Information Base implemented?	7.4.2	M	Yes ☐
NPDU Forwarding	Is forwarding of ISO8473 PDUs implemented?	7.4.3.1, 7.4.3.3	M	Yes ☐
Receive Process	Are the basic receive process functions implemented?	7.4.4	M	Yes ☐

A.6 Protocol summary: ISO 10589 level 1 specific functions

Item	Functionality/Description	References	Status	Support
*L1IS	Are Level 1 IS-IS routing functions implemented?	12.1.3	M	Yes ☐
Maximum Area Addresses	What values of maximumAreaAddresses are supported by this implementation?	7.1.5, 7.2.11	L1IS:M	values=
Area IS Count	How many ISs can this system support in a single area?	12.2.4	L1IS:M	N =
L1 Manual ES Adjacency	Are the manual ES adjacencies implemented?	7.3.3.1	L1IS:M	Yes ☐

A.6.1 Level 1 subnetwork dependent functions

Item	Functionality/Description	References	Status	Support
L1 LAN IIH PDU	Are L1 LAN IIH PDUs correctly constructed and parsed?	9.5	C.3:M	N/A ☐ Yes ☐

C.3 if L1IS and LAN then M else –

A.6.2 Level 1 update process

Item	Functionality/Description	References	Status	Support
L1 LS PDU	Are L1 LS PDUs correctly constructed and parsed?	9.8	L1IS: M	Yes ☐
L1 CSN PDU	Are L1 CSN PDUs correctly constructed and parsed?	9.10	L1IS:M	Yes ☐
L1 PSN PDU	Are L1 PSN PDUs correctly constructed and parsed?	9.12	L1IS:M	Yes ☐

A.6.3 Level 1 decision process

Item	Functionality/Description	References	Status	Support
L1 Nearest L2 IS Identification	Is the identification of the nearest L2 IS implemented?	7.2.9.1	L1IS:M	Yes ☐
L1 Area Addresses Computation	Is the computation of area addresses implemented?	7.2.11	L1IS:M	Yes ☐

A.7 Protocol summary: ISO 10589 level 2 specific functions

Item	Functionality/Description	References	Status	Support
*L2IS	Are Level 2 IS-IS routing functions implemented?	12.1.4	O	Yes ☐ No ☐
IS Count	What is the total number of ISs that this L2 IS can support?	12.2.4	L2IS: M	N/A ☐ N=
L2IS Count	How many level 2 ISs does this implementation support?	12.2.4.1	L2IS:M	N/A ☐ N=
*RA Prefix	Are Reachable Address Prefixes supported on circuits?	8.1, 7.3.3.2	L2IS:O	N/A ☐ Yes ☐ No ☐
External Metrics	Are external metrics supported?	7.2.2, 7.2.12, 7.3.3.2	RA Prefix:M	N/A ☐ Yes ☐
*Partition	Is level 1 partition repair implemented?	7.2.10	L2IS:O	N/A ☐ Yes ☐ No ☐

A.7.1 Level 2 subnetwork dependent functions

Item	Functionality/Description	References	Status	Support
L2 LAN IIH PDU	Are L2 LAN IIH PDUs correctly constructed and parsed?	9.6	C.4/M	N/A ☐ Yes ☐
*8208 DA	Are ISO8208 Dynamic Assigned circuits implemented?	8.3	O.1	Yes ☐ No ☐
RA Adjacency Management	Are the reachable address adjacency management operations implemented?	8.3.2.2-8.3.5.6	8208 DA:O	N/A ☐ Yes ☐
Call Establishment Metric Increment	Are non-zero values of the callEstablishmentMetric-Increment supported?	8.3.5	8208 DA:O	N/A ☐ Yes ☐ No ☐
Reverse Path Cache	Is 8208 reverse path cache implemented?	8.3.3	8208 DA:O	N/A ☐ Yes ☐ No ☐

C.4 if L2IS and LAN then M else –

A.7.2 Level 2 update process

Item	Functionality/Description	References	Status	Support
L2 LS PDU	Are L2 LS PDUs correctly constructed and parsed?	9.9	L2:M	N/A ☐ Yes ☐
L2 CSN PDU	Are L2 CSN PDUs correctly constructed and parsed?	9.11	L2:M	N/A ☐ Yes ☐
L2 PSN PDU	Are L2 PSN PDUs correctly constructed and parsed?	9.13	L2:M	N/A ☐ Yes ☐

A.7.3 Level 2 decision process

Item	Functionality/Description	References	Status	Support
L2 Attached Flag	Is the setting of the attached flag implemented?	7.2.9.2	L2IS:M	N/A ☐ Yes ☐
L2 Partition DIS election	Is the election of the partition L2 DIS implemented?	7.2.10.2	Partition:M	N/A ☐ Yes ☐
L2 Partition Area Addresses Computation	Is the computation of L1 partition area addresses implemented/	7.2.10.3	Partition:M	N/A ☐ Yes ☐
L2 DIS Partition Repair	Is partition detection and repair via virtual L1 links implemented?	7.2.10.1	Partition:M	N/A ☐ Yes ☐

A.7.4 Level 2 forward/receive process

Item	Functionality/Description	References	Status	Support
L2 NPDU Encapsulation	Is the encapsulation of NPDUs implemented?	7.2.10.4, 7.4.3.2	Partition:M	N/A ☐ Yes ☐
L2 NPDU Decapsulation	Is the decapsulation of NPDUs implemented?	7.4.4	Partition:M	N/A ☐ Yes ☐
L2 Encapsulated NPDU Forwarding	Is the encapsulated NPDU forwarding implemented	7.2.10.5 7.4.3	L2IS:M	N/A ☐ Yes ☐

Annex B (informative)

Supporting technical material

B.1 Addressing and routing

In order to ensure the unambiguous identification of Network and Transport entities across the entire OSIE, some form of address administration is mandatory. ISO 8348 specifies a hierarchical structure for network addresses, with a number of top-level domains responsible for administering addresses on a world-wide basis. These address registration authorities in turn delegate to subauthorities the task of administering portions of the address space. There is a natural tendency to repeat this subdivision to a relatively fine level of granularity in order to ease the task of each sub-authority, and to assign responsibility for addresses to the most “localised” administrative body feasible. This results in (at least in theory) reduced costs of address administration and reduced danger of massive address duplication through administrative error. Furthermore, political factors come into play which require the creation of sub-authorities in order to give competing interests the impression of “hierarchical parity”. For example at the top level of the ISO geographic address space, every country is assigned an equally-sized portion of the address space even though some countries are small and might in practice never want to undertake administration of their own addresses. Other examples abound at lower levels of the hierarchy, where divisions of a corporation each wish to operate as an independent address assignment authority even though this is inefficient operationally and may waste monumental amounts of potential address space.

If network topologies and traffic matrices aligned naturally with the hierarchical organisation of address administration authorities, this profligate use of hierarchy would pose little problem, given the large size (20 octets) of the N-address space. Unfortunately, this is not usually the case, especially at higher levels of the hierarchy. Network topologies may cross address administration boundaries in many cases, for example:

- Multi-national Corporations with a backbone network that spans several countries
- Community-of-interest networks, such as academic or research networks, which span organisations and geographies
- Military networks, which follow treaty alignments rather than geographic or national administrations
- Corporate networks where divisions at times operate as part of a contractor’s network, such as with trade consortia or government procurements.

These kinds of networks also exhibit rich internal topologies and large scale (10^5 systems), which require sophisticated routing technology such as that provided by this International Standard. In order to deploy such networks effectively, a considerable amount of address space must be left over for assignment in a way which produces efficient routes without undue consumption of memory and bandwidth for routing overhead¹⁾.

Similarly important is the inter-connection of these networks via Inter-domain routing technology. If all of the assignment flexibility of the addressing scheme is exhausted in purely administrative hierarchy (at the high-order end of the address) and in Intra-Domain routing assignment (at the low end of the address) there may be little or no address space left to customise to the needs of inter-domain routing. The considerations for how addresses may be structured for the Intra- and Inter-domain cases are discussed in more detail in the following two clauses.

B.1.1 Address structure for intra-domain routing

The IS-IS Intra-domain routing protocol uses a “preferred” addressing scheme. There are a number of reasons the designers of this protocol chose to specify a single address structure, rather than leaving the matter entirely open to the address assignment authorities and the routing domain administrators:

- a) If one address structure is very common and known a priori, the forwarding functions can be made much faster;

¹⁾ In other words, hierarchical routing, with its natural effect on address assignment, is a mandatory requirement for such networks.

- b) If part of the address is known to be assigned locally to an end system, then the routing can be simpler, use less memory, and be potentially faster, by not having to discriminate based on that portion of the address.
- c) If part of the address can be designated as globally unique by itself (as opposed to only the entire address having this property) a number of benefits accrue:
 - 1) Errors in address administration causing duplicate addresses become much less likely
 - 2) Automatic and dynamic NSAP address assignment becomes feasible without global knowledge or synchronisation
 - 3) Routing on this part of the address can be made simple and fast, since no address collisions will occur in the forwarding database.
- d) If a part of the address can be reserved for assignment purely on the basis of topological efficiency (as opposed to political or address administration ease), hierarchical routing becomes much more memory and bandwidth efficient, since the addresses and the topology are in close correspondence.
- e) If an upper bound can be placed on the amount of address space consumed by the Intra-domain routing scheme, then the use of address space by Inter-domain routing can be made correspondingly more flexible.

The preferred address format of the Intra-domain IS-IS protocol achieves these goals by being structured into fields as follows shown in figure B.1 below:

The field marked **AFI and IDI** in the figure are precisely the IDP specified in ISO 8348. The field marked **DSP** is treated as logically containing two parts. The high-order part is that portion of the DSP from ISO 8348 whose structure, assignment, and meaning are *not* specified or constrained by the Intra-domain IS-IS routing protocol. However, the design presumes that the routing domain administrator has at least some flexibility in assigning a portion of the DSP field. The purpose and usage of the fields specified by the Intra-domain IS-IS routing protocol is explained in the following paragraphs.

B.1.1.1 Area address

Since the Intra-domain IS-IS protocol is customised for operation with ISO 8473, all addresses are specified to use the preferred binary encoding of ISO 8348.

B.1.1.2 The *Selector (SEL)* Field

The SEL field is intended for two purposes. Its main use is to allow for multiple higher-layer entities in End systems (such as multiple transport entities) for those systems which need this capability. This allows up to 256 NSAPs in a single End system. The advantage of reserving this field exclusively for local system administration the Intra-domain routing functions need not store routing information about, nor even look at this field. If each individual NSAP were represented explicitly in routing tables, the size of these tables would grow with the number of NSAPs, rather than with the number of End systems. Since Intra-domain routing routes to *systems*, explicit recording of each NSAP brings no efficiency benefit and potentially consumes large amounts of memory in the Intermediate systems.




Figure B.1 - Preferred Address Format

A second use for the **SEL** field is in Intermediate systems. Certain IS–IS functions require that PDUs be encapsulated and sent to the Network Entity in an Intermediate system rather than to an NSAP and upward to a Transport entity. An example of this is the Partition Repair function of this International Standard. In order to use a level 2 path as if it were a single subnetwork in a level 1 area, PDUs are encapsulated and addressed to an IS on the other side of the partition¹⁾. By reserving certain values of the **SEL** field in Intermediate systems for direct addressing of Intermediate system Network entities, the normal addressing and relaying functions of other Intermediate systems can be transparently used for such purposes.

B.1.1.3 The *identifier* (id) field

The ID field is a “flat”, large identifier space for identifying OSI systems. The purpose of this field is to allow very fast, simple routing to a large (but not unconstrained) number of End systems in a routing domain. The Intra-Domain IS–IS protocol uses this field for routing within a *area*. While this field is only required to be unambiguous within a single area, if the values are chosen to be globally unambiguous the Intra-domain IS–IS design can exploit this fact in the following ways.

First, a certain amount of parallelism can be obtained during relaying. An IS can be simultaneously processing the **ID** field along with other fields. If the **ID** is found in the forwarding table, the IS can initiate forwarding while checking to make sure that the other fields have the expected value. Conversely, if the **ID** is not found the IS can assume that either the addressed NSAP is unreachable or exists only in some other area or routing domain. In the case where the **ID** is not globally unique, the forwarding table can indicate this fact and relaying delayed until the entire address is analysed and the route looked up.

Second, a considerable savings can be obtained in manual address administration for all systems in the routing domain. If the **ID** is chosen from the ISO 8802 48-bit address space (represented as a sequence of 6 octets as specified in ISO/IEC 10039), the **ID** is known to be globally unique. Furthermore, since LAN systems conforming to ISO 8802 often have their 48-bit MAC address stored in ROM locally, each system can be guaranteed to have a globally unambiguous NET and NSAP(s) without centralised address administration at the area level²⁾. This not only eliminates administrative overhead, but also drastically reduces the possibility of duplicate NSAP addresses, which are illegal, difficult to diagnose, and often extremely difficult to isolate.

An alternative to a large, flat space for the lowest level of routing would be to hierarchically subdivide this field to allow more levels of routing within a single routing domain. The designers of the Intra-domain IS–IS protocol considered that this would lead to an inferior routing architecture, since:

- a) The cost of memory in the ISs was sufficiently reasonable that large (e.g. 10 4 systems) areas were quite feasible, thus requiring at least 2 octets per level to address
- b) Two levels of routing within a routing domain were sufficient (allowing domains of 10 6 –10 7 systems) because it was unlikely that a single organisation would wish to operate and manage a routing domain much larger than that.
- c) Administrative boundaries often become the dominant concern once routing domains reach a certain size.
- d) The additional burdens and potential for error in manual address assignment were deemed serious enough to permit the use of a large, flat space.

B.2 Use of the area address field in intra-domain routing

Use of a portion of the **DSP** field provides for hierarchical routing within a routing domain. A value is assigned to a set of ISs in order to group the ISs into a single area for the usual benefits of hierarchical routing:

- a) Limiting the size of routing tables in the ISs;
- b) conserving bandwidth by hierarchical summarisation of routing information;
- c) designating portions of the network which are to have optimal routing within themselves; and
- d) moderate firewalling of portions of the routing domain from failures in other portions.

¹⁾ This is a gross oversimplification for the purpose of illustrating the need for the **SEL** field. See 7.2.10

²⁾ Note, however, that the use of the ISO 8802 addresses does not avoid the necessity to run ISO 9542 or to maintain tables mapping NSAP addresses to MAC (i.e. SNPA) addresses on the ISO 8802 subnetwork. This is because there is no guarantee that a particular MAC address is always enabled (the LAN controller may be turned off) or that a system has only a single MAC address.

It is important to note that the assignment of high-order **DSP** values is intended to provide the routing domain administrator with a mechanism to optimise the routing within a large routing domain. The Intra-domain IS-IS designers did *not* intend the **DSP** to be entirely consumed by many levels of address registration authority. Reserving the assignment of a portion of the **DSP** field to the routing domain administrator also allows the administrator to start with a single assigned **Area Address** and run the routing domain as a single area. As the routing domain grows, the routing domain administrator can then add areas without the need to go back to the address administration authority for further assignments. Areas can be added and reassigned within the routing domain without involving the external address administration authority.

A useful field to reserve as part of the **DSP** would be 2 octets, permitting up to 65 536 areas in a routing domain. This is viewed as a reasonable compromise between routing domain size and address space consumption. The field may be specified as flat for the same reasons that the ID field may be flat.

B.2.1 Addressing considerations for inter-domain routing

It is in the Inter-domain arena where the goals of routing efficiency and administrative independence collide most strongly. Although the OSI Routing Framework explicitly gives priority to considerations of autonomy and firewalls over efficiency, it must be *feasible* to construct an Inter-Domain topology that both produces isolable domains and relays data at acceptable cost. Since no routing information is exchanged across domain boundaries with static routing, the practicality of a given Inter-domain topology is essentially determined by the size of the routing tables that are present at the boundary ISs. If these tables become too large, the memory needed to store them, the processing needed to search them, and the bandwidth needed to transmit them within the routing domain all combine to disallow certain forms of interconnection.

Inter-domain routing primarily computes routes to other routing domains¹⁾. If there is no correspondence between the address registration hierarchy and the organisation of routing domains (and their interconnection) then the task of static table maintenance quickly becomes a nightmare, since each and every routing domain in the OSIE would need a table entry potentially at every boundary IS of every other routing domain. Luckily, there is some reason to believe that a natural correspondence exists, since at least at the global level the address registration authorities fall within certain topological regions. For example, most of the routing domains which obtained their Area Addresses from a hierarchy of French authorities are likely to reside in France and be more strongly connected with other routing domains in France than with routing domains in other countries.

There are enough exceptions to this rule, however, to be a cause for concern. The scenarios cited in B.1 all exist today and may be expected to remain common for the foreseeable future. Consider as a practical case the High Energy Physics Network (HEPnet), which contains some 17 000 End systems, and an unknown number of intermediate systems²⁾. This network operates as a single routing domain in order to provide a known set of services to a known community of users, and is funded and cost-justified on this basis. This network is international in scope (at least 10 countries in North America, Europe, and the Far East) and yet its topology does not map well onto existing national boundaries. Connectivity is richer between CERN and FERMILab, for example than between many points within the U.S.

More importantly, this network has rich connectivity with a number of other networks, including the PDNs of the various countries, the NSFnet in the U.S., the international ESnet (Energy Sciences Network), the general research Internet, and military networks in the U.S. and elsewhere. None of these other networks shares a logical part of the NSAP address hierarchy with HEPnet³⁾. If the only method of routing from the HEPnet to these other networks was to place each within one and only one of the existing registration authorities, and to build static tables showing these relationships, the tables would clearly grow as $O(n^2)$.

It seems therefore, that some means must be available to assign addresses in a way that captures the Inter-Domain topology, and which coexists cleanly with both the administrative needs of the registration authorities, and the algorithms employed by both the Intra- and Inter-domain routing protocols. As alluded to in an earlier clause, it seems prudent to leave some portion of the address space (most likely from the high order part of the **DSP**) sufficiently undefined and flexible that various Inter-domain topologies may be efficiently constructed.

¹⁾ This International Standard also uses static Inter-domain tables for routing to individual End systems across dynamically assigned circuits, and also to End systems whose addresses do not conform to the address construction rules.

²⁾ The number of ISs is hard to estimate since some ISs and links are in fact shared with other networks, such as the similarly organised NASA Space Physics network, or "SPAN"

³⁾ It is conceivable that ISO would sanction such networks by assigning a top-level IDI from the ISO non-geographic AFI, but this is unlikely and would only exacerbate the problem if many such networks were assigned top-level registrations.

Annex C (informative)

Implementation guidelines and examples

C.1 Routing databases

Each database contains records as defined in the following sub-clauses. The following datatypes are defined.

```

FROM CommonMgmt IMPORT NAddress,
AddressPrefix, BinaryAbsoluteTime;
lspID = ARRAY [0..iDLength-1] OF Octet;
systemID = ARRAY [0..iDLength] OF Octet;
octetTimeStamp = BinaryAbsoluteTime;

```

C.1.1 Level 1 link state database

This database is kept by Level 1 and Level 2 Intermediate Systems, and consists of the latest Level 1 Link State PDUs from each Intermediate System (or pseudonode) in the area. The Level 1 Link State PDU lists Level 1 links to the Intermediate System that originally generated the Link State PDU.

RECORD

```

adr: lspID;
(* ID of LSP originator *)
type: (Level1IntermediateSystem,
      AttachedLevel2IntermediateSystem,
      UnattachedLevel2IntermediateSystem);
seqnum: [0..SequenceModulus - 1];
LSPage: [0..MaxAge];
(*Remaining Lifetime *)

expirationTime: TimeStamp;
(*Time at which LSP age became zero (see 7.3.16.4). *)
SRMflags: ARRAY[1..(maximumCircuits +
maximumVirtualAdjacencies)] OF BOOLEAN;
(*Indicates this LSP to be sent on this circuit. Note that level 2 Intermediate systems may send level 1 LSPs to
other partitions (if any exist). Only one level 2 Intermediate system per partition does this. For level 1 Intermediate
Systems the array is just maximumCircuits long. *)
SSNflags: ARRAY[1..maximumCircuits +
maximumVirtualAdjacencies] OF BOOLEAN;
(*Indicates that information about this LSP shall be included in the next partial sequence number PDU transmitted
on this circuit. *)
POINTER TO LSP;
(*The received LSP *)
END;

```

C.1.2 Level 2 link state database

This database is kept by Level 2 Intermediate Systems, and consists of the latest Level 2 Link State PDUs from each Level 2 Intermediate System (or pseudonode) in the domain. The Level 2 Link State PDU lists Level 2 links to the Intermediate System that originally generated the Link State PDU.

RECORD

```

adr: lspID;
(* iDLength + 2 octet ID of LSP originator *)
type: (AttachedLevel2IntermediateSystem,
      UnattachedLevel2IntermediateSystem);

```

```

seqnum: [0..SequenceModulus – 1];
LSPage: [0..MaxAge];
(*Remaining Lifetime *)
expirationTime: TimeStamp;
(*Time at which LSP age became zero (see
7.3.16.4). *)
SRMflags: ARRAY[1..(maximumCircuits)] OF
BOOLEAN;
(*Indicates this LSP to be sent on this circuit. *)
SSNflags: ARRAY[1..maximumCircuits] OF
BOOLEAN;
(*Indicates that information about this LSP must be included in the next partial sequence number PDU transmitted
on this circuit. *)
POINTER TO LSP; (*The received LSP *)
END;

```

C.1.3 Adjacency database

This database is kept by all systems. Its purpose is to keep track of neighbours. For Intermediate systems, the adjacency database comprises a database with an entry for each:

- Adjacency on a Point-to-point circuit.
- Broadcast Intermediate System Adjacency.

NOTE 78 Both a Level 1 and a Level 2 adjacency can exist between the same pair of systems

- Broadcast End system Adjacency.
- potential SVC on a DED circuit (maximumSVCAdjacencies for a DA circuit, or 1 for a Static circuit).
- Virtual Link Adjacency.

Each entry contains the parameters in Clause 11 for the Adjacency managed object. It also contains the variable used to store the remaining holding time for each Adjacency IDEntry, as defined below.

```

IDEntry = RECORD
ID: systemID;
(* The idLength octet System ID of a neighbour End system extracted from the SOURCE ADDRESS field of its
ESH PDUs. *)
entryRemainingTime: Unsigned [1..65535]
(* The remaining holding time in seconds for this entry. This value is not accessible to system management. An
implementation may choose to implement the timer rules without an explicit remainingTime being maintained. For
example by the use of asynchronous timers. It is present here in order to permit a consistent description of the timer
rules. *)
END

```

Circuit Database

This database is kept by all systems. Its purpose is to keep information about a circuit. It comprises an **ARRAY**[1..maximumCircuits].

Each entry contains the parameters in Clause 11 for a linkage managed object. It also contains the remainingHelloTime (WordUnsigned [1..65535] seconds) variable for the Circuit. This variable **not** accessible to system management. An implementation may choose to implement the timer rules without an explicit remainingHelloTime being maintained. For example by the use of asynchronous timers. It is present here in order to permit a consistent description of the timer rules. Additionally, for Circuits of type “X.25 Static Outgoing” or “X.25 DA”, it contains the recallCount (Unsigned[0..255]) variable for the Circuit. This variable is **not** accessible to system management. It used to keep track of recall attempts.

C.1.4 Level 1 shortest paths database

This database is kept by Level 1 and Level 2 Intermediate Systems (unless each circuit is “Level 2 Only”). It is computed by the Level 1 Decision Process, using the Level 1 Link State Database. The Level 1 Forwarding Database is a subset of this database.

RECORD

adr: systemId; (*idLength octet ID of destination system *)
 cost: [1..MaxPathMetric];
 (*Cost of best path to destination system *)
 adjacencies: ARRAY[1..maximumPathSplits]
OF POINTER TO Adjacency;
 (*Pointer to adjacency for forwarding to system adr *)
END;

C.1.5 Level 2 shortest paths database

This database is kept by Level 2 Intermediate Systems. It is computed by the Level 2 Decision Process, using the Level 2 Link State Database. The Level 2 Forwarding Database is a subset of this database.

RECORD

adr: AddressPrefix;
 (*destination prefix *)
 cost: [1..MaxPathMetric];
 (*Cost of best path to destination prefix *)
 adjacencies: ARRAY[1..maximumPathSplits]
OF POINTER TO Adjacency;
 (*Pointer to adjacency for forwarding to prefix adr *)
END;

C.1.6 Level 1 forwarding database

This database is kept by Level 1 and Level 2 Intermediate Systems (unless each circuit is “Level 2 Only”). It is used to determine where to forward a data NPDU with destination within this system’s area. It is also used to determine how to reach a Level 2 Intermediate System within the area, for data PDUs with destinations outside this system’s area.

RECORD

adr:systemId;
 (*idLength octet ID of destination system. Destination “0” is special, meaning “nearest level 2 Intermediate system” *)
 splits: [0..maximumPathSplits];
 (* Number of valid output adj’s for reaching adr (0 indicates it “is unreachable”) *)
 nextHop: ARRAY[1..maximumPathSplits] **OF**
POINTER TO adjacency;
 (*Pointer to adjacency for forwarding to destination system *)
END;

C.1.7 Level 2 forwarding database

This database is kept by Level 2 Intermediate systems. It is used to determine where to forward a data NPDU with destination outside this system’s area.

RECORD

adr: AddressPrefix;
 (*address of destination area. *)
 splits: [0..maximumPathSplits];
 (*Number of valid output adj’s for reaching adr (0 indicates it is unreachable) *)
 nextHop: ARRAY[1..maximumPathSplits] **OF**
POINTER TO adjacency;
 (*Pointer to adjacency for forwarding to destination area. *)
END;

C.2 SPF algorithm for computing equal cost paths

An algorithm invented by Dijkstra (see references) known as *shortest path first (SPF)*, is used as the basis for the route calculation. It has a computational complexity of the square of the number of nodes, which can be decreased to the number of links in the domain times the log of the number of nodes for sparse networks (networks which are not highly connected).

A number of additional optimisations are possible:

- a) If the routing metric is defined over a small finite field (as in this International Standard), the factor of $\log n$ may be removed by using data structures which maintain a separate list of systems for each value of the metric rather than sorting the systems by logical distance.
- b) Updates can be performed incrementally without requiring a complete recalculation. However, a full update must be done periodically to recover from data corruption, and studies suggest that with a very small number of link changes (perhaps two) the expected computation complexity of the incremental update exceeds the complete recalculation. Thus, this International Standard specifies the algorithm only for the full update.
- c) If only End system LSP information has changed, it is not necessary to recompute the entire Dijkstra tree for the IS. If the proper data structures exist, End Systems may be attached and detached as leaves of the tree and their forwarding information base entries altered as appropriate.

The original SPF algorithm does not support load splitting over multiple paths. The algorithm in this International Standard does permit load splitting by identifying a set of equal cost paths to each destination rather than a single least cost path.

C.2.1 Databases

PATHS – This represents an acyclic directed graph of shortest paths from the system S performing the calculation. It is stored as a set of triples of the form $\langle N, d(N), \{Adj(N)\} \rangle$, where:

N is a system Identifier. In the level 1 algorithm, N is a $(idLength+1)$ octet ID. For a non-pseudonode it is the $idLength$ octet system ID, with a 0 appended octet. For a pseudonode it is a true $(idLength+1)$ octet quantity, comprised of the $idLength$ octet Designated Intermediate System ID and the extra octet assigned by the Designated Intermediate System. In the level 2 algorithm it is either a $(idLength+1)$ octet Intermediate System or pseudonode ID (as in the level 1 algorithm), or it is a variable length address prefix (which will always be a leaf, i.e. “End system”, in PATHS).

$d(N)$ is N 's distance from S (i.e. the total metric value from N to S).

$\{Adj(N)\}$ is a set of valid adjacencies that S may use for forwarding to N . When a system is placed on PATHS, the path(s) designated by its position in the graph is guaranteed to be a shortest path.

TENT – This is a list of triples of the form $\langle N, d(N), \{Adj(N)\} \rangle$, where N , $d(N)$ and $\{Adj(N)\}$ are as defined above for PATHS.

TENT can intuitively be thought of as a tentative placement of a system in PATHS. In other words, the triple $\langle N, x, \{A\} \rangle$ in TENT means that if N were placed in PATHS, $d(N)$ would be x , but N cannot be placed on PATHS until it is guaranteed that no path shorter than x exists.

The triple $\langle N, x, \{A, B\} \rangle$ in TENT means that if N were placed in PATHS, $d(N)$ would be x via either adjacency A or B

NOTE 79 As described above, (see 7.2.6), it is suggested that the implementation keep the database TENT as a set of lists of triples of the form $\langle *, Dist, * \rangle$, for each possible distance $Dist$. In addition it is necessary to be able to process those systems which are pseudonodes before any non-pseudonodes at the same distance $Dist$.

C.2.2 Use of metrics in the SPF calculation

Internal metrics are not comparable to external metrics. Therefore, the cost of the path from N to S for external routes (routes to destinations outside of the routing domain) may include both internal and external metrics. The cost of the path from N to S (called $d(N)$ below in database PATHS) may therefore be maintained as a two-dimensional vector quantity (specifying internal and external metric values). In incrementing $d(N)$ by 1, if the internal metric value is less than the maximum value $MaxPathMetric$, then the internal metric value is incremented by one and the external metric value left unchanged; if the

internal metric value is equal to the maximum value **MaxPathMetric**, then the internal metric value is set to 0 and the external metric value is incremented by 1. Note that this can be implemented in a straightforward manner by maintaining the external metric as the high order bits of the distance.

NOTE 80 In the code of the algorithm below, the current path length is held in a variable **tentlength**. This variable is a two-dimensional quantity **tentlength**=(internal,external) and is used for comparing the current path length with $d(N)$ as described above.

C.2.3 Overview of the algorithm

The basic algorithm, which builds **PATHS** from scratch, starts out by putting the system doing the computation on **PATHS** (no shorter path to **SELF** can possibly exist). **TENT** is then pre-loaded from the local adjacency database.

Note that a system is not placed in **PATHS** unless no shorter path to that system exists. When a system N is placed in **PATHS**, the path to each neighbour M of N , through N , is examined, as the path to N plus the link from N to M . If $\langle M, *, * \rangle$ is in **PATHS**, this new path will be longer, and thus ignored.

If $\langle M, *, * \rangle$ is in **TENT**, and the new path is shorter, the old entry is removed from **TENT** and the new path is placed in **TENT**. If the new path is the same length as the one in **TENT**, then the set of potential adjacencies $\{adj(M)\}$ is set to the union of the old set (in **TENT**) and the new set $\{adj(N)\}$. If M is not in **TENT**, then the path is added to **TENT**.

Next the algorithm finds the triple $\langle N, x, \{Adj(N)\} \rangle$ in **TENT**, with minimal x .

NOTE 81 This is done efficiently because of the optimisation described above. When the list of triples for distance **Dist** is exhausted, the algorithm then increments **Dist** until it finds a list with a triple of the form $\langle *, \text{Dist}, * \rangle$.

N is placed in **PATHS**. We know that no path to N can be shorter than x at this point because all paths through systems already in **PATHS** have already been considered, and paths through systems in **TENT** will have to be greater than x because x is minimal in **TENT**.

When **TENT** is empty, **PATHS** is complete.

C.2.4 Algorithm

The Decision Process Algorithm must be run once for each supported routing metric. A Level 1 Intermediate System runs the algorithm using the Level 1 LSP database to compute Level 1 paths. In addition a Level 2 Intermediate System runs the algorithm using the Level 2 LSP database to compute Level 2 paths.

If this system is a Level 2 Intermediate System which supports the partition repair optional function the Decision Process algorithm for computing Level 1 paths must be run twice for the default metric. The first execution is done to determine which of the area's **manualAreaAddresses** are reachable in this partition, and elect a Partition Designated Level 2 Intermediate System for the partition. The Partition Designated Level 2 Intermediate System will determine if the area is partitioned and will create virtual Level 1 links to the other Partition Designated Level 2 Intermediate Systems in the area in order to repair the Level 1 partition. This is further described in 7.2.10.

C.2.5 Step 0: Initialise **TENT** and **PATHS** to empty. Initialise **tentlength** to (0,0).

(**tentlength** is the pathlength of elements in **TENT** we are examining.)

- Add $\langle \text{SELF}, 0, W \rangle$ to **PATHS**, where W is a special value indicating traffic to **SELF** is passed up to Transport (rather than forwarded).
- Now pre-load **TENT** with the local adjacency database. (Each entry made to **TENT** must be marked as being either an End system or an Intermediate System to enable the check at the end of Step 2 to be made correctly.) For each adjacency $Adj(N)$, (including Manual Adjacencies, or for Level 2 enabled Reachable Addresses) on enabled circuits, to system N of **SELF** in state "Up", compute

$d(N)$ = cost of the parent circuit of the adjacency (N), obtained from metric k , where k = one of *default metric, delay metric, monetary metric, error metric*.

$Adj(N)$ = the adjacency number of the adjacency to N .

- c) If a triple $\langle N, x, \{Adj(M)\} \rangle$ is in TENT, then:
If $x = d(N)$, then $Adj(M) \leftarrow \{Adj(M)\} \cup Adj(N)$.
- d) If there are now more adjacencies in $\{Adj(M)\}$ than **maximumPathSplits**, then remove excess adjacencies as described in 7.2.7.
- e) If $x < d(N)$, do nothing.
- f) If $x > d(N)$, remove $\langle N, x, \{Adj(M)\} \rangle$ from TENT and add the triple $\langle N, d(N), Adj(N) \rangle$.
- g) If no triple $\langle N, x, \{Adj(M)\} \rangle$ is in TENT, then add $\langle N, d(N), Adj(N) \rangle$ to TENT.
- h) Now add any systems to which the local Intermediate system does not have adjacencies, but which are mentioned in neighbouring pseudonode LSPs. The adjacency for such systems is set to that of the Designated Intermediate System.
- i) For all broadcast circuits in state "On", find the LSP with LSP number zero and with the first (**idLength**+1) octets of LSPID equal to the L n CircuitID for that circuit (i.e. pseudonode LSP for that circuit). If it is present, for all the neighbours N reported in all the LSPs of this pseudonode which do not exist in TENT add an entry $\langle N, d(N), Adj(N) \rangle$ to TENT, where

$$d(N) = \text{metric } k \text{ of the circuit.}$$

$$Adj(N) = \text{the adjacency number of the adjacency to the DR.}$$
- j) Go to **Step 2**.

C.2.6 Step 1: Examine the zeroth Link State PDU of P, the system just placed on PATHS (i.e. the Link State PDU with the same first (**idLength**+1) octets of LSPID as P, and LSP number zero).

- a) If this LSP is present, and the **LSP Database Overload** bit is clear, then for each LSP of P (i.e. all the Link State PDUs with the same first (**idLength**+1) octets of LSPID as P, irrespective of the value of LSP number) compute

$$dist(P, N) = d(P) + \text{metric } k(P, N).$$

for each neighbour N (both Intermediate System and End system) of the system P. If the **LSP Database Overload** bit is set, only consider the End system neighbours of the system P. $d(P)$ is the second element of the triple $\langle P, d(P), \{Adj(P)\} \rangle$ and $\text{metric } k(P, N)$ is the cost of the link from P to N as reported in P's Link State PDU
- b) If $dist(P, N) > \text{MaxPathMetric}$, then do nothing.
- c) If $\langle N, d(N), \{Adj(N)\} \rangle$ is in PATHS, then do nothing.
- d) If a triple $\langle N, x, \{Adj(N)\} \rangle$ is in TENT, then:
 - 1) If $x = dist(P, N)$, then $Adj(N) \leftarrow \{Adj(N)\} \cup Adj(P)$.
 - 2) If there are now more adjacencies in $\{Adj(N)\}$ than **maximumPathSplits**, then remove excess adjacencies, as described in 7.2.7.
 - 3) If $x < dist(P, N)$, do nothing.

NOTE 82 $d(N)$ must be less than $dist(P, N)$, or else N would not have been put into PATHS. An additional sanity check may be done here to ensure $d(N)$ is in fact less than $dist(P, N)$

- 4) If $x > dist(P, N)$, remove $\langle N, x, \{Adj(N)\} \rangle$ from TENT and add $\langle N, dist(P, N), \{Adj(P)\} \rangle$.
- e) If no triple $\langle N, x, \{Adj(N)\} \rangle$ is in TENT, then add $\langle N, dist(P, N), \{P\} \rangle$ to TENT.

C.2.7 Step 2: If TENT is empty, stop, else:

a) Find the element $\langle P, x, \{Adj(P)\} \rangle$, with minimal x as follows:

- 1) If an element $\langle *, tentlength, * \rangle$ remains in TENT in the list for **tentlength**, choose that element. If there are more than one elements in the list for **tentlength**, choose one of the elements (if any) for a system which is a pseudonode in preference to one for a non-pseudonode. If there are no more elements in the list for **tentlength**, increment **tentlength** and repeat Step 2.
- 2) Remove $\langle P, tentlength, \{Adj(P)\} \rangle$ from TENT.
- 3) Add $\langle P, d(P), \{Adj(P)\} \rangle$ to PATHS.
- 4) If this is the Level 2 Decision Process running, and the system just added to PATHS listed itself as Partition Designated Level 2 Intermediate system, then additionally add $\langle AREA.P, d(P), \{adj(P)\} \rangle$ to PATHS, where AREA.P is the Network Entity Title of the other end of the Virtual Link, obtained by taking the first AREA listed in P's Level 2 LSP and appending P's ID.
- 5) If the system just added to PATHS was an End system, go to Step 2, Else go to Step 1.

NOTE 83 In the Level 2 context, the "End systems" are the set of Reachable Address Prefixes and the set of **areaAddresses** with zero cost.

C.3 Forwarding process**C.3.1 Example pseudo-code for the forwarding procedure described in 7.4.3**

This procedure chooses, from the Level 1 forwarding database – if **level** is **level1**, or from the Level 2 forwarding database – if **level** is **level2**, an adjacency on which to forward PDUs for destination **dest**. A pointer to the adjacency is returned in **adj**, and the procedure returns the value "True". If no suitable adjacency exists the procedure returns the value "False", in which case a call should be made to Drop("Destination Address Unreachable", **octetNumber**).

If queue length values are available to the forwarding process, the minimal queue length of all candidate circuits is chosen, otherwise, they are used in round robin fashion.

```

PROCEDURE Forward(
    level: (level1, level2),
    dest: NetworkLayerAddress,
    VAR adj: POINTER TO
    adjacency) : BOOLEAN
VAR
    adjArray: ARRAY OF
    ForwardingDatabaseRecords;
    temp, index, minQueue: CARDINAL;
BEGIN
    (*Set adjArray to appropriate database*)
    IF level = level1 THEN
        adjArray := level1ForwardingDatabase
    ELSE
        adjArray := level2ForwardingDatabase
    END;
    (*Perform appropriate hashing function to obtain an index into the database*)
    IF Hash(level, dest, index) THEN
        IF adjArray[index].splits > 0 THEN
            (*Find minimum queue size for all equal cost paths*)
            minQueue := MaxUnsigned;
            temp := adjArray[index].lastChosen + 1;
            (*start off after last time*)
            FOR i := 1 TO adjArray[index].splits DO
                (*for all equal cost paths to dest*)
                IF temp > adjArray[index].splits THEN
                    (*after end of valid entries, wrap to first*)

```



```
    temp := 1
  ELSE
    temp := temp + 1
  END;
  IF
    QueueSize(adjArray[index].nextHop[temp]) < minQueue THEN
      minQueue :=
        QueueSize(adjArray[index].nextHop[temp]);
      adj :=
        adjArray[index].nextHop[temp];
      adjArray[index].lastChosen := temp;
    END;
    Forward := true
  END;
  ELSE
    Forward := false (*There must be at least one valid output adjacency *)
  END
ELSE
  Forward := false (*Hash returned destination unknown *)
END
END forward;.
```

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:2002

Annex D (informative)

Congestion control and avoidance

D.1 Congestion control

The transmit management subroutine handles congestion control. Transmit management consists of the following components:

Square root limiter. Reduces buffer occupancy time per PDU by using a square root limiter algorithm. The square root limiter also queues PDUs for an output circuit, and prevents buffer deadlock by discarding PDUs when the buffer pool is exhausted. D.1.1 specifies the Square Root Limiter Process.

Originating PDU limiter. Limits originating NPDU traffic when necessary to ensure that transit NPDUs are not rejected. An originating NPDU is an NPDU resulting from an NSDU from the Transport at this ES. A transit NPDU is an NPDU from another system to be relayed to another destination ES.

Flusher. Flushes PDUs queued for an adjacency that has gone down.

Information for higher layer (Transport) congestion control procedures is provided by the setting of the “congestion experienced” bit in the forwarded NPDUs.

D.1.1 Square root limiter

The square root limiter discards a data NPDU by calling the ISO 8473 discard PDU function with the reason “PDU Discarded due to Congestion” when the number of data NPDUs on the circuit output queue exceeds the discard threshold, U_d . U_d is given as follows:

where:

$$U_d = \left\lceil \frac{N_b}{\sqrt{N_c}} \right\rceil$$

N_b = Number of Routeing Layer buffers (maximumBuffers) for all output circuits.

N_c = Number of active output circuits (i.e. Circuits in state “On”).

The output queue is a queue of buffers containing data NPDUs which have been output to that circuit by the forwarding process, and which have not yet been transmitted by the circuit. It does **not** include NPDUs which are held by the data link layer for the purpose of retransmission.

Where a data NPDU is to be fragmented by this Intermediate system over this circuit, each fragment shall occupy a separate buffer and shall be counted as such in the queue length. If the addition of all the buffers required for the fragmentation of a single input data NPDU would cause the discard threshold for that queue to be exceeded, it is recommended that all those fragments (including those which could be added without causing the threshold to be exceeded) be discarded.

D.1.2 Originating PDU limiter

The originating PDU limiter first distinguishes between *originating* NPDUs and *transit* NPDUs. It then imposes a limit on the number of buffers that originating NPDUs can occupy on a per circuit basis. In times of heavy load, originating NPDUs may be rejected while transit NPDUs continue to be routed. This is done because originating NPDUs have a relatively short wait, whereas transit NPDUs, if rejected, have a long wait – a transport retransmission period.

The originating PDU limiter accepts as input:

- An NSDU received from Transport Layer
- A transmit complete signal from the circuit for an ISO 8473 Data PDU.

The originating PDU limiter produces the following as output:

- PDU accepted
- PDU rejected
- Modifications to originating PDU counter

There is a counter, N , and an originating PDU limit, `originatingQueueLimit`, for each active output circuit. Each N is initialised to 0. The `originatingQueueLimit` is set by management to the number of buffers necessary to prevent the circuit from idling.

D.1.3 Flusher

The flusher ensures that no NPDU is queued on a circuit whose state is not ON, or on a non-existent adjacency, or one whose state is not Up.

D.2 Congestion avoidance

D.2.1 Buffer management

The Forwarding Process supplies and manages the buffers necessary for relaying. PDUs shall be discarded if buffer thresholds are exceeded. If the average queue length on the input circuit **or** the forwarding processor **or** the output circuit exceeds `QueueThreshold`, the “congestion experienced” bit shall be set in the QoS maintenance option of the forwarded data PDU (provided the QoS maintenance option is present).

IECNORM.COM : Click to view the full PDF of ISO/IEC 10589:2002

Annex E (normative)

Syntax imported from ISO 10165-5 (SC6 GMI)

--E.1 Generic managed object class definitions

MODULE "Rec. X.723 | ISO/IEC 10165-5"

--E.1.1 Communication information record

-- The communicationsInformationRecord object class is used to define
 -- the information stored in a log as a result of receiving event reports
 -- with a communication information event type. The semantics of the
 -- object class, namely its attributes and behaviour, are derived from
 -- the communicationsInformation notification.
 --

communicationsInformationRecord MANAGED OBJECT CLASS
 DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992":eventLogRecord;
 CHARACTERIZED BY communicationsInformationRecordP1 PACKAGE
 BEHAVIOUR communicationsInformationRecordB1 BEHAVIOUR
 DEFINED AS
 !Log record class for communicationsInformation events.!
 ;
 ATTRIBUTES
 informationType GET;
 ;;
 CONDITIONAL PACKAGES
 informationDataPackage PACKAGE
 ATTRIBUTES
 informationData GET;
 REGISTERED AS {CommonMgt.poi informationDataPackage(1001)};
 PRESENT IF !The informationData parameter is present in the
 communicationsInformation event report corresponding to
 the instance of communicationsInformationRecord.!
 REGISTERED AS {CommonMgt.moi communicationsInformationRecord(1101)};

--E.1.2 Communications entity

-- Summary of changes to 10165-5:
 --
 -- 1. Added CCITT part to document references.
 -- 2. Gave plural labels to set-valued attributes.
 -- 3. Fixed typo in label communicationsEntityId.
 -- 4. Did not register mandatory package.
 --

communicationsEntity MANAGED OBJECT CLASS
 DERIVED FROM "Rec. X.721 | ISO/IEC 10165-2 : 1992":top;
 CHARACTERIZED BY communicationsEntityP1 PACKAGE
 BEHAVIOUR
 communicationsEntityB1 BEHAVIOUR
 DEFINED AS
 !A communications entity supports the disabled and enabled
 values of the operationalState attribute as described in
 CCITT Rec. X.731 | ISO/IEC 10164-2 as follows: