Stable Implementation Agreements for Open Systems Interconnection Protocols: Part 3 - Network Layer

Output from the March 1994 Open Systems Environment Implementors' Workshop (OIW)

SIG Chair: **Fred Burg, AT&T**
SIG Editor: **Brenda Gray, NIST**

Foreword

This part of the Stable Implementation Agreements was prepared by the Lower Layers Special Interest Group (LLSIG) of the Open Systems Environment Implementors' Workshop (OIW). See Part 1 - Workshop Policies and Procedures of the "Draft Working Implementation Agreements Document" for the charter.

Text in this part has been approved by the Plenary of the above-mentioned Workshop. ~~This part replaces the previously existing chapter on this subject.~~

Annex A is for information only.

Future changes and additions to this version of these Implementor Agreements will be published as change pages. Deleted and replaced text will be shown as strikeout. New and replacement text will be shown as shaded.

Table of Contents

Part 3 - Network Layer March 1994 (Stable)

List of Figures

List of Tables

Part 3 - Network Layer

0 Introduction

This part presents agreements for providing the OSI network service. Also contained here are agreements on network layer addressing and routing.

Scope

These agreements cover both connectionless-mode and connection-mode network services.

Normative References

CCITT

[**1**] Recommendation X.213 (Blue Book, 1988), *Network Service Definition for Open Systems Interconnection for CCITT Applications*.

ISO

[**2**] ISO 8348, *Information processing systems - Data communications - Network service definition*.

[**3**] ISO 8348 Addendum 1, *Information processing systems - Data communications - Network service definition - Addendum 1: Connectionless-mode transmission.*

[**4**] ISO 8348 Addendum 2, *Information processing systems - Data communications - Network service definition - Addendum 2: Network layer addressing*.

[**5**] ISO 8473, *Information processing systems - Data communications - Protocol for providing the connectionless-mode network service*.

[**6**] ISO 8648, *Information processing systems - Open systems interconnection - Internal organization of the Network Layer*.

[**7**] ISO 8878, *Information processing systems - Data communications - Use of X.25 to provide the OSI connection-mode network service*.

[**8**] ISO 8881, *Information processing systems - Data communications - Use of the X.25 packet level protocol in local area networks*.

[**9**] ISO 9542, *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 service (ISO 8473)*.

[**10**] ISO/IEC 9574, *Information technology - Telecommunications and information*

exchange between systems - Provision of the OSI connection-mode network service by packet mode terminal equipment connected to an integrated services digital network (ISDN).

[**11**] ISO/IEC TR 9577, *Information technology - Telecommunications and information exchange between systems - Protocol identification in the network layer*.

[**12**] ISO/IEC TR 10029, *Information technology - Telecommunications and information exchange between systems - Operation of an X.25 interworking unit*.

[**13**] ISO/IEC 10030, *Information processing systems - Telecommunications and information exchange between systems - End system routeing information exchange protocol for use in conjunction with ISO 8878*.

[**14**] ISO/IEC 10589, *Information technology - Telecommunications and information exchange between systems - Intermediate system to intermediate system intra-domain routeing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473).*

[**15**] ISO/IEC DIS 11577, *Information Technology - Telecommunications and Information Exchange Between Systems - Network Layer Security Protocol*

Status

This version of the agreements was completed in December 1993.

Errata

This clause may contain "Defect Report" and resolutions material, and the versions of implementor agreements to which this material applies.

The following defects are being progressed in ISO:

ISO 9542, defect 1, Parts 1-13;

ISO 9542, defect 2;

ISO 8473, defects 1 through 11, and technical corrigendum 1;

ISO/IEC 10589, defect 1;

ISO/IEC 10589, defect 2.

Connectionless-Mode Network Service (CLNS)

ISO 8473

NOTE - Defect reports upon the base standard have been issued. See clause 4 of this part for further information.

Subsets of the Protocol

Agreements on subsets of the protocol are as follows:

Implementations will not transmit PDUs encoded using the inactive subset. Received PDUs encoded using the inactive subset will be discarded;

The non-segmenting subset will not be used. Implementations will not generate data PDUs without a segmentation part. However, implementations will receive and correctly process PDUs which do not contain the segmentation part.

Mandatory Functions of ISO 8473

Agreements on Mandatory Functions of ISO 8473 are as follows:

The lifetime parameter shall be used as specified in clause 6.4 of ISO 8473. The parameter shall have an initial value of at least three times the network span or three times the maximum transit delay (in units of 500 ms), whichever is greater;

The reassembly timer for an initial PDU at the reassembly point shall be no greater than the largest value of all lifetime parameters contained in all derived PDUs;

The use/non-use of checksums shall be capable of being configured. The default setting shall be non-use;

If the implementation supports the generation of an ER PDU, the system shall insert in the destination address field of the ER PDU the contents of the source address field of the PDU that generated the error;

For the purposes of relaying and routing, a protocol entity need not verify the correctness of ISO 8348/Add. 2 semantics carried in NPAI of received PDUs.

Optional Functions of ISO 8473

Agreements on Optional Functions of ISO 8473 are as follows:

The Security parameter is not defined by these Agreements. Implementations shall not transmit the parameter except where defined by bilateral agreements;

Partial and complete source routing will not be supported;[1]

Partial record of route will be supported by Intermediate systems;

ISO 8473 will be followed with respect to QOS;

For systems implementing the congestion notification function, the following applies:

A Globally Unique QOS Maintenance parameter shall be included in all PDU

[1] A defect exists with the Partial Source Routing option which can cause PDUs to loop in the network until their lifetime expires.

originated by End Systems. As specified in ISO 8473, the initial value of the Congestion Experienced flag (CE flag) within the Globally Unique QOS Maintenance Parameter shall be set by the originating End System to zero. All other flags within the Globally Unique QOS Maintenance Parameter shall be set based on the specific local needs of the originating End System;

Intermediate systems not implementing queue length averaging shall leave the CE flag in the same state as it was received. In particular, no intermediate system (IS) shall ever clear (set to zero) the CE flag. All intermediate systems shall monitor all incoming and outgoing queues and compute average queue lengths as shown by example in figure 1. The averaging is done from the beginning of the previous cycle to the current time. A cycle begins at the instant of the first NSDU arrival after an idle period;

An IS should set the CE flag in all NSDUs forwarded on a queue which has an average queue length greater than one;

The queue length averaging algorithm computes the average queue length over two cycles, where the two cycles are:

> the "previous cycle", which is the interval from when the IS becomes busy, until it becomes idle and the idle ends (indicated by the instant the first packet arrives to the idle IS);
>
> the "current cycle", which is the interval from the end of the idle interval to the current time instant when the average queue length is computed;

An embodiment of the averaging algorithm is shown in figure 1;

Refer to the Working Implementation Agreements document for additional optional functions.

The algorithm makes use of the following variables:

t = Current time
t_i = time of i^{th} arrival or departure event
q_i = number of packets in the system after the event
T_0 = time at the beginning of the previous cycle
T_1 = time at the beginning of the current cycle

The algorithm consists of three components:

1. Queue Length Update: Beginning with $q_0 = 0$,
   If the i^{th} event is an arrival event, $q_i = q_{i-1}+1$
   If the i^{th} event is a departure event, $q_i = q_{i-1}-1$

2. Queue Area (integral) update:

$$\text{Area of the previous cycle} = \sum_{t_i \varepsilon \{T_0,T_1)} q_{i-1}(t_i-t_{i-1})$$

$$\text{Area of the current cycle} = \sum q_{i-1}(t_i-t_{i-1})$$

$t_i \varepsilon \{T_1, t)$

3. Average Queue Length Update:

Average Queue length over the two cycles

$$= \frac{\text{Area of the two cycles}}{\text{Time of the two cycles}} = \frac{\text{Area of the two cycles}}{t - T_0}$$

Figure 1 - Queue length averaging algorithm

Provision of CLNS over Local Area Networks (LANS)

When providing CLNS over a LAN subnetwork, the following shall apply:

The definition of CLNS shall be as specified in ISO 8348/Add. 1;

The protocol used to provide CLNS shall be ISO 8473 with agreements as specified in 5.1;

The necessary subnetwork dependent convergence function shall be as defined in ISO 8473 - clause 8.4.2, "SNDCF used with ISO 8802/2 sub-networks."

Provision of CLNS over X.25 Subnetworks

When providing CLNS over X.25 subnetworks, the following shall apply:

The definition of CLNS shall be as specified in ISO 8348/Add. 1;

The protocol used to provide CLNS shall be ISO 8473 with agreements as specified in 5.1;

The necessary subnetwork dependent convergence function shall be as defined in ISO 8473 - clause 8.4.3, "SNDCF used with ISO 8208 subnetworks for operation over X.25 subnetworks," and the default throughput class shall be used if this facility is available;

The X.25 PLP shall be as specified in part 2 clause 6.3.

Provision of CLNS over ISDN

When providing CLNS over an ISDN, the following shall apply:

The definition of CLNS shall be as specified in ISO 8348/Add. 1;

The protocol used to provide CLNS shall be ISO 8473 with agreements as specified in 5.1;

The necessary Subnetwork Dependent Convergence function shall be as defined in:

ISO 8473 for operation of CLNP over X.25 with agreements as specified in 5.3;

ISO 9574 for control of the B and D channels;

The X.25 PLP shall be as specified in part 2, clause 6.3;

The agreements for the ISDN-related protocols are specified in part 2, clause 7.

NOTE - The stated scope of ISO 9574 does not explicitly cover the operation of CLNP over an ISDN. However, the procedure identified for operating X.25 in conjunction with I.451 is still applicable. The procedures in ISO 9574 that correspond to 8878 are not utilized when providing CLNS.

Provision of CLNS over Point-to-Point Links

Refer to the Working Implementation Agreements document.

Connection-Mode Network Service (CONS)

The following agreements concern provision of the connection-mode Network Service.

Mandatory Method of Providing CONS

General

Independent of the subnetwork type (of part 2), when providing the CONS using X.25-1984, the following shall apply as described below:

The definition of the CONS is as specified in ISO 8348, Network Service Definition;

The mapping of the elements of the CONS to the elements of the X.25 Packet Layer Protocol (PLP) is as specified in 6.3.1;

The general procedures and formats of the X.25 PLP are as specified in ISO 8208, X.25 Packet Layer Protocol for Data Terminal Equipment.

X.25 WAN

No provisions additional to those in 6.1.1 apply in an X.25 WAN.

LANs

When providing the CONS in a Local Area Network, the following aspects of ISO 8881, in addition to the documents listed in 6.1.1, shall apply:

Clauses 1-6 and 9-11 for LLC Type 1 operation, including the additional nonstandard default packet size listed in Clause 6.3, Note 2.

NOTE - Operation of ISO 8208 in conjunction with LLC Type 2 requires agreement on LLC Type 2 procedures.

ISDN

When providing the CONS in an ISDN, the considerations for control of a B and D channel in ISO 9574, in addition to those provided in 6.1.1, shall apply.

Additional Option: Provision of CONS over X.25 1980 Subnetworks

When providing CONS over an X.25 1980 subnetwork, the following shall apply:

The definition of the CONS is as specified in ISO 8348, Network Service Definition;

The subnetwork dependent convergence protocol required to provide CONS shall be as specified in ISO 8878 Annex A, and referred to as the Alternative Procedures for Network Connection Establishment and Release, with agreements as defined in 6.3.2.

Agreements on Protocols

ISO 8878

ISO 8878 Clauses 1-11 shall apply with the following exception:

Where the ISO 8208 diagnostic codes are not provided, all Cause/Diagnostic code combinations can be mapped to the Originator/Reason code of "Undefined."

Subnetwork Dependent Convergence Protocol (ISO 8878/Annex A)

The Receipt Confirmation service will not be provided, so the corresponding protocol elements need not be implemented.

The Expedited Data service will not be provided, so the corresponding protocol elements need not be implemented.

Interworking

Interworking between subnetworks whose End Systems use ISO 8208 to provide the CONS as specified in 6.1 shall be performed as specified in ISO TR 10029. That is, an Intermediate System connecting two such subnetworks shall operate ISO 8208 on both subnetworks and shall relay information from one subnetwork to the other as described in ISO TR 10029.

Addressing

NSAP address formats supported will conform to Addendum 2 of ISO 8348 as follows:

NSAP address formats will have a hierarchical structure. This will reduce the size of routing tables;

If used in the Domain Specific Part (DSP), an NSAP selector shall be the least significant component in the hierarchy, and shall be encoded as the last octet of the DSP. The NSAP selector shall not be used to perform routing; it is simply intended to identify the network service user at the destination end system. For those implementations using an NSAP selector, there shall be one and only one selector for each NSAP within the end system. All NSAP addresses identifying a given NSAP will use the same NSAP selector value.

c) In routing environments in which systems support NSAP addresses containing selectors as specified in b), the corresponding Network Entity Titles shall have the same format with the NSAP selector set to zero.

d) End Systems and Intermediate Systems operating in routing domains that employ the ISO 10589 Intradomain Routing Protocol shall meet the NSAP/NET addressing requirements specified in ISO 10589 (clause 7.1) and clause 8.3 of these agreements.

NOTE - This may be incompatible with systems implemented according to previous versions of these agreements.

Routing

The basic principles of Network Layer routing are defined in the OSI Routing Framework ISO/IEC TR 9575. These principles state that:

The global OSI environment will consist of a number of Administrative Domains, An Administrative Domain consists of a collection of End Systems (ESs) and Intermediate Systems (ISs), and subnetworks operated by a single organization or Administrative Authority. The Administrative Authority is responsible for: the organization of ESs and ISs into Routing Domains; the assignments of NSAP and SNPA addresses; the policies that govern resource usage; the policies that govern the information that is collected and disseminated both internally and externally to the Administrative Domain; and the establishment of subdomains and the corresponding delegation of responsibilities;

A Routing Domain is a set of ESs and ISs which operate according to the same routing procedures and which is wholly contained within a single Administrative Domain. An Administrative Authority may delegate to the entity responsible for a Routing Domain the responsibilities to further structure and assign NSAP and SNPA addresses. The hierarchical decomposition of Routing Domains into subdomains may greatly reduce the resources required in the maintenance, computation, and storage of routing information;

The OSI routing problem, and consequently OSI routing protocols, has been decomposed into three distinct classes:

End System (ES) to Intermediate System (IS) routing within a single subnetwork;

IS to IS routing within a single routing domain (Intra-domain);

IS to IS routing between routing domains (Inter-domain).

ISO 9542 End System to Intermediate System Routing

NOTE - Defect reports upon the base standard have been issued. See clause 4 of this part for further information.

For use in conjunction with ISO 8473, ISO 9542 shall be used to provide the routing exchange protocol.

Additionally, a management mechanism capable of adding and deleting entries into the Routing Information Base (RIB) is recommended. When using the management mechanism to add an entry, there should be no holding timer, and the entry should be write protected from alteration by the ES-IS protocol. This mechanism enables routing table entries to be made which are static in nature.

The agreements below apply to the use of ISO 9542:

Implementors shall support any valid NSAP format. For the purposes of the protocol, NSAP addresses are treated simply as octet strings;

For LANs, implementors shall support both Configuration Information and Route Redirection Information; no subsets are permitted. For X.25 subnetworks, Route Redirection Information shall be supported;

All timer values shall be configurable;

Use or non-use of checksums shall be configurable. It is recommended not to use ISO 9542 checksums when originating PDUs;

The QOS, Security and Priority parameters should not be used for routing. For conformance, intermediate systems must transmit these parameters in RD PDUs if they are present in the data PDU which generated the redirect. However, end systems must ignore them in received RD PDUs;

If the configuration notification function described in 6.7 of the protocol specification is implemented, a mechanism shall be provided to enable/disable this function on broadcast networks. If supported in end systems listening to both ISHs and ESHs, this function shall only be invoked upon receipt of an ISH. Alternate mechanisms for ISs and ESs are described in 8.1.1 and 8.1.2.

For LANs, this protocol employs the same LSAP as ISO 8473;

The encoding of the BSNPA address follows the syntax rules for the data link being used. On a LAN, for example, it is a 48-bit MAC address encoded as specified in clause 12.2.1.4 of ISO 10039. On X.25 subnetworks, it is a DTE address, each digit being binary coded in a semi-octet, and, if there are an odd number of digits, an additional semi-octet set to the value 1111 shall be added at the end;

The multicast addresses corresponding to "All Intermediate Systems on the Network" (ALL_ISN) and "All End Systems on the Network" (ALL_ESN) shall default to the following on IEEE802.3 and IEEE802.4 subnetworks:

ALL_ESN = 09-00-2B-00-00-04, ALL_ISN = 09-00-2B-00-00-05;

It is understood that the hexadecimal octets shown are transmitted onto the medium from left most octet to right most octet. Within each hexadecimal octet the least significant bit is transmitted first;

when operating on a specific IEEE 802.5 subnetwork, all ESs and ISs shall use exclusively either Functional Addresses or Group Addresses for the operation of ISO 9542. It is strongly recommended that Group Addressing be used where possible.

For IEEE 802.5 LANs in which Group Addressing is supported by all ESs and ISs, the multicast addresses corresponding to "All Intermediate Systems on the Network" (ALL_ISN) and "All End Systems on the Network" (ALL_ESN) shall be as follows:

ALL-ESN = 09-00-2B-00-00-04, ALL_ISN = 09-00 2B-00-00-05;

For IEEE 802.5 LANs in which Functional Addressing must be used, the ISO 9542 multicast shall be as follows:

1) ALL-ESN = 03-00-00-00-02-00, ALL_ISN = 03-00-00-00-01-00.

Editor's Note - When transmitted onto the medium, the above addresses would be represented as follows:

1) ALL_ESN = C0-00-00-00-40-00, ALL_ISN = C0-00-00-00-80-00.

The Error Report flag shall be set to zero (0) for NPDUs sent as a result of invoking the QUERY Configuration Function.

ISO 8473 PDUs multicast as a result of the Query Configuration function shall use the Network Layer Protocol ID (NLPID) assigned to ISO 8473.

An ISO 8473 PDU received as a result of another ES having performed the Query Configuration function shall be processed as follows:

If the ISO 8473 PDU is addressed to one of the NSAPs present in the ES, the End System shall process the PDU according to the applicable clauses of ISO 8473 and invoke the Configuration Response Function (clause 6.6 of ISO 9542);

If the ISO 8473 PDU is not addressed to one of the NSAPs present in the ES, the End System shall discard the PDU without generating an ISO 8473 Error Report;

For purposes of address matching and SNPA extraction, the first octet of the option parameter value of an address (clause 7.4.5) or SNPA Mask (clause 7.4.6) shall be aligned with the first octet (AFI) of the encoded trial NSAP Address.

to enable the operation of the Configuration Information subset by End Systems attached to an X.25 subnetwork, an End System may optionally be configured with the SNPA addresses of Intermediate Systems on the subnetwork.

The following items represent proposed solutions to defects in ISO 9542. These solutions are being progressed as defect reports to ISO 9542. These items will be deleted when the corresponding defect report is approved:

An End System may choose to ignore an RD PDU received for a destination to which the ES has not sent traffic for some period of time. An ES must record redirection information only for those other systems with which it is in active communication;

A holding time value of zero is permitted. When configuration and/or redirection information with a zero holding time is received, prior information shall be replaced, thus causing the system to set its holding timer to zero and discard the corresponding information;

If one or more ISs suggested an ESCT, the minimum of the non-zero suggested values replaces the current value of the ES's CT.

Alternative Configuration Mechanism - IS Actions

An alternative mechanism for achieving rapid configuration which is scaleable to large broadcast networks is described below. This mechanism makes use of the Suggested ES Configuration Timer. Implementation of this mechanism is optional.

When an Intermediate system wants to quickly acquire the End system configuration (for example, when a broadcast circuit is enabled on the IS or the topology changes because of a failure of a bridge or repeater), it initiates a "poll" of the End system configuration by performing the following actions:

Delay a random interval between 0 and PollESHelloRate seconds. (This is to avoid synchronization with other ISs which have detected a change.);

In order to rapidly time out any End systems which are no longer present on the broadcast circuit (for example, after a LAN partition), reset the entryRemainingTime in the Routing Information Base (RIB) for all End systems on this circuit to the value: **(ISHelloTimer + PollESHelloRate) * HoldingMultiplier** or the existing value whichever is lowest. Where ISHelloTimer is the Intermediate system's configuration timer, HoldingMultiplier is a predefined number (for example, 2) which multiplied by ISHelloTimer gives the value for the Holding Time field of IS Hellos;

Then transmit HoldingMultiplier IS Hellos with a Suggested ES Configuration Timer value of PollESHelloRate seconds with an interval of ISHelloTimer seconds between each and setting the Holding Time field to ISHelloTimer * HoldingMultiplier;

Then start sending IS Hellos with a Suggested ES Configuration Timer of DefaultESHelloRate seconds (where DefaultESHelloRate is larger than PollESHelloRate).

Alternate Configuration Mechanism - ES Actions

An End system maintains for each circuit a list (CTList) which has HoldingMultiplier elements each of which stores a received value of the Suggested ES Configuration Timer. The function SaveCT(t) adds the value t as the first element of CTList and discards the last element. The function MinCT delivers the minimum value in CTList. When the circuit is enabled all the elements of CTList are initialized to PollESHelloRate.

An End system also maintains for each circuit the variables currentSuggestedHelloTimer and its associated lifetime currentSuggestedHelloTimerLifetime. These are both initialized to PollESHelloRate.

When the circuit is enabled the Configuration Timer is started by setting the entryRemainingTime to random (PollESHelloRate).

On Configuration Timer expiry the following actions are performed:

SaveCT(currentSuggestedHelloTimer);

Transmit an ES Hello with Holding Time field set to MinCT * HoldingMultiplier;

Set entryRemainingTime to MinCT - random(MinCT * 0.25). (The random element ensures that End systems do not become synchronized.)

When an End system receives an IS Hello which contains a Suggested ES Configuration Timer, it is processed as follows (where suggestedESCT is the value contained in the option):

If suggestedESCT is less than or equal to currentSuggestedHelloTimer then set curentSuggestedHelloTimerLifetime to the value of the Holding Time field of the IS Hello;

If suggestedESCT is less than currentSuggestedHelloTimer then set currentSuggestedHelloTimer to suggestedESCT and reset entryRemainingTime to the smaller of its current value and random(currentSuggestedHelloTimer * 0.75).

When the currentSuggestedHelloTimerLifetime expires, set the currentSuggestedHelloTimer to DefaultESHelloTimer.

ISO 10030 End System to Intermediate System Routing

The protocol used to provide End System to Intermediate System routing in support of the CONS (refer to 3.6) shall be ISO 10030.

The following agreements apply to the use of ISO 10030:

A management mechanism capable of adding and deleting entries in the Routing Information Base (RIB) of both SNAREs and End Systems is recommended. When using the management mechanism to add an entry it should not be timed out, and the entry should be write protected from alteration by the ISO 10030 protocol.

The multicast addresses corresponding to "All CONS End Systems" and "All CONS SNAREs" shall default to the following on IEEE 802.3 and IEEE 802.4 subnetworks:

All CONS End Systems = 01-80-C2-00-00-16

All CONS SNAREs = 01-80-C2-00-00-17

Intra-Domain Intermediate Systems to Intermediate Systems Routing

Static Intra-Domain Routing

Intermediate systems shall provide mechanisms to create and update the required Routing Information Base (RIB).

Dynamic Intra-Domain Routing

NOTE - Defect reports upon the base standard have been issued. See clause 4 of this part for further information.

The protocol used to provide Intermediate System to Intermediate System routing in support of the CLNS (refer to clause 3.5) among systems in a single routing domain shall be ISO 10589.

The following agreements apply to the use of ISO 10589:

A management mechanism capable of configuring the Identifier, Characteristic, and Status attributes of the managed objects of clause 11 shall be provided;

The implementation shall support a system identifier (ID) length of 6 octets and shall use this value as a default;

When operating on IEEE 802.5 (i.e., token ring) LANs, the group addresses specified in ISO/IEC 10589, clause 8.4.8, table 9 shall be used.

Inter-Domain Intermediate Systems to Intermediate Systems Routing

Static Inter-Domain Routing

Intermediate Systems shall provide ~~management~~ mechanisms to create and update the required Routing Information Base (RIB).

Dynamic Inter-Domain Routing

The protocol used to provide the exchange of inter-domain routing information among intermediate systems in support of the CLNS (refer to clause 3.5) shall be ISO/IEC DIS 10747 (IDRP).

The following agreements apply to the support and use of IDRP:

A management mechanism capable of configuring the attributes and reporting the notifications defined in the management information of clause 11 shall be provided;

The authentication mechanisms of clause 7.9 shall be supported.

IDRP provides a wide range of protocol mechanisms capable of supporting the exchange of "path attributes" associated with routing information and allows implementations to support various "policies" for controlling the selection of routes and the subsequent distribution of routing information.

The set of path attributes and policies supported largely determines the complexity of an implementation and its ability to interoperate with other IDRP systems in support of a given domain's specific routing policies.

In order to promote the ability of independent IDRP systems to be deployed and interoperate to effect routing policies, implementations shall support one of the following configuration subsets:

Minimum Implementation Subset (MIS)

Implementations supporting the MIS provide minimal capabilities to encode and effect routing policies.

The following path attributes must be supported:

RD_HOP_COUNT;

RD_PATH;

LOC_PREF;

EXT_INFO;

DIST_LIST_INCL;

DIST_LIST_EXCL;

An IDRP implementation supporting the MIS shall support the following routing policy mechanisms:

Default - all potential destination are acceptable;

List Controlled - distribution lists determine which destinations domains are acceptable;

For these policies, only a single "default" routing information base and forwarding information base is required. No distinguishing attributes are supported by the MIS.

Enhanced Implementation Subset (EIS):

Implementations supporting the EIS shall provide all the capabilties of the MIS and the additional capabilities described below.

The following path attributes must be supported:

MULTI_EXIT_DISC;

NEXT_HOP;

HIERARCHICAL_RECORDING.

An EIS implementation shall support routing domain confederations and corresponding routing information aggregation capabilities.

An IDRP implementation supporting the EIS shall support the routing policy mechanisms of the MIS with the following additions:

Intermediate Source - distribution lists control the acceptance and/or propagation of routing information based upon the identify of the adjacent systems providing the information;

General Path - the ability to control the acceptance and/or propagation of routing information based upon the contents of the PATH attributes.

For these policies at least a single "default" routing information base and forwarding information base is required. Additional distinguishing attributes (and corresponding information bases) may be supported.

Procedures for OSI Network Service/Protocol Identification

General

The Protocol Identifiers specified in ISO TR 9577 ("Protocol Identification in the OSI Network Layer") provide a basis from which OSI systems (both end systems and intermediate systems) may derive a set of procedures for indicating which OSI protocols are used in a particular instance of communication. As such, these procedures are only concerned with Initial Protocol Identifiers (IPIs) and Subsequent Protocol Identifiers (SPIs) that identify OSI protocols and pertain to the following types of systems:

systems providing/supporting only CONS (using ISO 8208/8878);

systems providing/supporting only CLNS (using ISO 8473);

systems providing/supporting both CONS and CLNS.

From this set of definitions, the following possibilities for success (S) or failure (F) of an instance of communication can be defined, as shown in the table below:

Table 1 - End Systems Communications

Originating	Destination End System Type		
End System Type	A	B	C
A	S	F	S
B	F	S	S
C	S	S	S

Processing of Protocol Identifiers

The usage of Protocol Identifiers in Network Protocol Data Units (NPDUs) depends on several factors:

the OSI Network Service to be provided;

the protocol to be used in providing this service;

the role the protocol is to be used in (per the Internal Organization of the Network Layer);

the type of subnetwork to which the system is connected.

Originating NPDUs

The use of a particular OSI Network Service depends on the capabilities of both the origination and destination end systems. It is not the intent of this clause to provide guidelines on how to make this choice except for simple obvious criteria; rather, it is intended only to provide guidance on how to convey this choice to the destination system.

Where a priori knowledge exists in the originating end system about the capabilities (with respect to OSI Network Services available) of the destination end system, it should be used. This may result in no communication if the two end systems involved only provide Network Services of different types. A selection is required in cases where both end systems provide both types of network services; this selection is conveyed by the use of the IPI and SPI (but the selection process is an implementation matter). Alternatively, where a priori knowledge does not exist, then the selection of a service to use in an instance of communication depends solely on the capabilities of the originating end system as described below:

If only CONS-related protocols (e.g., ISO 8208) are available, then this should be used and the Protocol Identifiers specified so as to reflect the chosen protocol(s) and service;

If only CLNS-related protocols (e.g., ISO 8473) are available, then this should be used and the Protocol Identifiers specified so as to reflect the chosen protocol(s) and service;

If both services are available, then other criteria are used in deciding which to use in an instance of communication.

NOTE - The choice of OSI Network Service to be used in an instance of communication is reflected in the Network Service primitives issued by the Network Service user.

Once a selection of Network Service has been made, the use of particular protocols depend on, for example, the subnetwork to which the originating End System is attached. Some specific cases are given in Annex A of ISO TR 9577. Another case involves use of the Protocol for Providing the Connectionless Network Service directly over the Data Link Service, as given in ISO 8473 (e.g., in a LAN). In this case, the IPI indicates ISO 8473.

Destination System Processing

A system receiving an NPDU must first be concerned with the protocol identified by the IPI. Valid values are given in table 2 of ISO TR 9577. If the protocol is recognized as one supported by the system, further processing of the protocol is performed according to the rules of that protocol. If not, an error is recognized and may be conveyed to the originating peer entity. With respect to ISO 8208 and ISO 8473, the following would apply for such error conditions:

For ISO 8208, the condition is classified as an "invalid General Format Identifier," for which a DIAGNOSTIC packet may be returned. If DIAGNOSTIC packets are not used by the system, the NPDU is discarded without any further action;

For ISO 8473, the NPDU is discarded without any further action.

Given acceptance of the protocol identified by the IPI, the system must also determine the acceptability of the subsequent protocols and OSI Network Service being requested. Use of ISO 8473 implies CLNS; however, use of ISO 8208 can imply either CONS or CLNS, as identified by the SPI. In the case of ISO 8208, therefore, further processing is needed to determine the acceptability of the requested protocol/service. If these are not acceptable (e.g., not supported by the system), the call should be cleared with a diagnostic code of "Connection Rejection - unrecognizable protocol identifier in user data" (decimal 249).

NOTE - In ISO 8208, a call may be refused for reasons other than non-support of the requested OSI Network Service.

Further Processing in Originating End System

Further processing on receipt of an NPDU in response to an initial attempt to communicate may be necessary/useful to determine the success of such an attempt.

For ISO 8473, when used directly over the Data Link Service, the success or failure of an attempt to communicate may not be visible/obvious within the Network Layer. On the other hand, use of ISO 8473 over ISO 8208 may provide, via the diagnostic code in a received CLEAR INDICATION packet, an indication of failure to communicate (e.g., the remote system does not support CLNS).

When using ISO 8208 to provide the CONS, the diagnostic code in a received CLEAR INDICATION packet may provide the necessary indication of why a call was refused.In cases where an ISO 8208 call is refused with diagnostic #249, it would not be desirable to re-attempt such calls with the exact same set of parameters; however, how the originating system ensures this is a local matter.

In cases where an originating system is capable of supporting both OSI Network Services, it may wish to re-attempt communications using the other mode of Network Service than that initially attempted.

Applicable Protocol Identifiers

The protocol identifiers applicable to these agreements are given in table 2 and table 3.

Table 2 - IPI Values

Bit Pattern 8 7 6 5 4 3 2 1	Protocol
0 0 0 0 1 0 0 0	CCITT I.451/Q.931
1 0 0 0 0 0 0 1	ISO 8473 (excluding the inactive subset)
1 0 0 0 0 0 1 0	ISO 9542
1 0 0 0 0 0 1 1	ISO/IEC 10589
1 0 0 0 0 1 1 0	ISO/IEC 11577
x x 0 1 x x x x	ISO 8208/CCITT X.25-Modulo 8
x x 1 0 x x x x	ISO 8208/CCITT X.25-Modulo 128
0 0 1 1 x x x x	ISO 8208/CCITT X.25-GFI Extension

Table 3 - SPI Values

Bit Pattern[1] 8 7 6 5 4 3 2 1	Protocol
0 0 0 0 0 0 0 0 thru 0 0 1 1 1 1 1 1	ISO 8073 ADD1/CCITT X.224 See table 4.1
1 0 0 0 0 0 0 1	ISO 8473

1 0 0 0 0 0 1 1	ISO/IEC 10589
1 0 0 0 0 1 0 0	ISO 8878/Annex A
1 0 0 0 0 1 1 0	ISO/IEC 11577

NOTES

1 A null SPI value (e.g., no Call User Data Field in an ISO 8208/CCITT X.25 Call Request/Incoming Call packet) shall indicate ISO 8073/CCITT X.224.

When using ISO 8208, values other than one of those listed in table 3 are outside the scope of these agreements.

Migration Considerations

This clause considers problems arising from evolving OSI standards and implementations based on earlier versions of OSI standards.

Until there is widespread availability of 1984 X.25 service, it will be necessary for X.400 systems to use those existing packet-switched public data networks which offer only pre-1984 X.25 service. While 1980 X.25 does not provide the CONS as defined by ISO 8348, there is no implication of non-conformance to these Agreements resulting therefrom for systems using 1980 X.25 to interchange data at the Network Layer, provided they conform in all other respects.

This is an exception to the Agreements for providing the OSI Network Service, granted temporarily for practical reasons. This exception will be removed when it is deemed to be no longer necessary, in the judgement of the Workshop. While this provision is in effect, it provides an alternative method of using 1980 X.25 to the provisions of 6.2.

Use of Priority

Refer to the Working Implementation Agreements document.

Introduction

Refer to the Working Implementation Agreements document.

Overview

Refer to the Working Implementation Agreements document.

Security

ISO/IEC DIS 11577 Network Layer Security Protocol (NLSP)

ISO/IEC DIS 11577 describes both a connction oriented and connectionless security protocol that can be used in conjunction with OSI Network Layer Protocols. Before secure communication can be accomplished, a security association (in band or out of band) shall have been established with agreement on all attributes associated with this association.

Managed objects are not yet specified by this standard and therefore the security domain/administrative authority shall determine the procedures and policies that govern this information with other security information.

All mandatory functions are supported by these implementation agreements.

Services

If access control service is selected and the label mechanism is used then integrity shall also be selected.

Mechanisms

To optimize efficiency and assist in the interoperability of secure implementations, it is useful to specify which mechanisms and algorithms apply. This specification shall allow implementations to know the exact encapsulation format used including what fields are required, their length, and order. A set of applicable profiles (mechanisms and algorithms) shall be specified within the Implementation Agreements to insure this efficient interoperability.

Protocol Data Unit

Although the standard has the option of all type-length-value (tlv) fields being in any order, for efficiency, the encapsulation format depicted in the standard shall be used. If the tlv fields are not in order, undefined (type field has not been allocated a value in the NLSP standard), or the PDU fails one of the NLSP Security checks, the secure encapsulated PDU should be discarded. The reporting of this situation is a local matter. If shared knowledge of this event is required, a possible technique would be to use the system management to report the error.

The Security Association-Identification field should be no more than twenty octets.

Functional Security Sequence Ordering

If Access control is implemented using labels, the label function is first applied followed by the integrity function. If confidentiality has also been selected, then that function is perfomed after the integrity function.

If integrity and confidentiality have been selected, then the integrity function is performed before the confidentiality function.

Conformance

Refer to the Working Implementation Agreements document.

Annex (informative)

Bibliography

CCITT Recommendation X.223 - 1988, Use of X.25 to Provide the OSI Connection-mode Network Service for CCITT Applications.

FIPS 100, Interface Between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) for Operation with Packet-Switched Data Communications Networks, NTIS, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161.

ISO/IEC 8880-1, Information Processing Systems - Data Communications - Protocol Combinations to Provide and Support the OSI Network Service - Part 1: General Principles.

ISO/IEC 8880-2, Information Processing Systems - Data Communications - Protocol Combinations to Provide and Support the OSI Network Service - Part 2: Provision and Support of the Connection-mode Network Service.

ISO/IEC 8880-3, Information Processing Systems - Data Communications - Protocol Combinations to Provide and Support the OSI Network Service - Part 3: Provision and Support of the Connectionless-mode Network Service.

ISO/IEC TR 9575, Information Technology - Telecommunications and Information Exchange Between Systems - OSI Routing Framework.