american national standards institute, inc.
1430 broadway, new york, n.y. 10018
(212) 354-3300

X3S3.3/91- 226

ISO/IEC JTC 1/SC6 N
1991-09-03

ISO
INTERNATIONAL ORGANIZATION FOR STANDARDIZATION
ORGANISATION INTERNATIONALE DE NORMALISATION

ISO/IEC JTC 1/SC 6

TELECOMMUNICATIONS AND INFORMATION EXCHANGE BETWEEN SYSTEMS

Secretariat: USA (ANSI)

TITLE: Approach for Providing OSI Connectionless-mode Network and Connectionless-mode Transport Multicast

SOURCE: D. Marlow (NSWC)

PROJECT: Enhanced Transport Mechanisms

ASSIGNED TO:

Note: Attached is an approach for providing multicast capabilities to the OSI Connectionless protocols at the Transport and Network Layers. This approach is based on the techniques developed by Steve Deering and described in RFC 1112.

This is a proposed contribution for the Ad Hoc SC6 meeting on Enhanced Transport Mechanisms scheduled for February 1992 in Paris. U.S. discussion on this input is planned for the ANSI X3S3.3 committee's September 1991 meeting in Boulder, CO.

While Annex E is not a candidate for international standardization, it is offered as supporting information for the February 1992 Ad Hoc SC6 Meeting.

1 Introduction

This paper provides an approach for providing connectionless multicast transfer capabilities to End Systems using the ISO OSI set of protocols. This paper identifies extensions to the “Protocol for providing the connectionless-mode network service” (CLNP, ISO 8473), “Addendum 2: Network layer addressing to the Network service definition” (ISO 8348 Addendum 2), “End system to Intermediate system routeing exchange protocol for use with the Protocol for providing the connectionless-mode network service” (ES-IS, ISO 9542), and “Protocol for providing the connectionless-mode transport service” (ISO 8602). While the issue of Intermediate System to Intermediate System routeing is not addressed at this time, considerations for future extensions to the OSI routeing protocols are used for justification of particular aspects of the approach contained here.

Initiating OSI multicast transfer capabilities in this way mirrors the present status of the Internet protocol suite where the means by which End Systems and Intermediate Systems interact in multicast transfer is complete but methods for Intermediate Systems for route determination and to exchange multicast PDUs among themselves are still in development. This approach meets current applications where multicast PDUs are needed to be exchanged on the same network or where PDUs traverse different networks using Static, Quasi-Static or Centralized routeing techniques. Techniques to extend OSI distributed adaptive routeing protocols to support multicast transfer are future work items.

2 Multicast concept of operation

The multicast concept of operation comes directly from RFC 1112 where connectionless PDU transfer is provided. OSI Connectionless multicasting is the transfer of a CLNP PDU to a set of zero or more End Systems identified by a single group NSAP address. The set of Network entities identified by such an address make up a particular group. Principal concepts are:

a) This proposal only considers multicast capability for the OSI connectionless protocols.

b) Multicast destinations are identified by group NSAP addresses.

c) Any End System in the interconnected OSI environment can listen on any group NSAP address.

d) Set of listeners may be dynamic (End Systems joining or leaving groups at any time) and unbounded.

e) Senders need not be listeners or know any information concerning listeners.

f) While no specific means of subnetwork multicasting is assumed by this approach some types of subnetworks (i.e. IEEE and FDDI LANs) have particular capabilities to facilitate multicast transfer.

Multicast capability is brought into the OSI Connectionless protocols via optional protocol extensions and addressing features where they are needed. Two basic assumptions are made: (1)Route determination and relaying of multicast PDU’s will be different than current practice and (2)Group Network addresses need to be “easily” distinguishable from non-group addresses. While the protocols governing multicast route determination are assumed to be different than non-multicast, the OSI structure for how the various protocols within the Network layer relate to each other remains intact.

The first assumption is justified by considering the consequences of using group addresses with the current Network layer protocols. With multicast transfer, End Systems will register via a group NSAP address to receive particular multicast PDUs and thus certain group NSAP addresses may be registered across the global network addressing domain. If Intermediate Systems accepted such group address registrations as normal End System registrations the result will be to interpret every registration of a group NSAP address as a separate SNPA of a very expansive multi-homed system. Such an interpretation of a group NSAP address is not consistent with the goals of multicast transfer. This leads to the need for separating the operations of Intermediate Systems to support multicast route determination from those that support the present functions. A primary goal is to not impact ESs and ISs which conform to the present protocol standards. Optional extensions provided here to ISO 9542 are used to achieve this objective.

The second assumption is justified by considering the present goals for the use of hierarchical structure of routeing domains for non-group NSAP address allocations. If group NSAPs addresses are permitted to be used throughout the global network addressing domain and these group NSAP addresses are assigned an address allocated to a particular area of a leaf routeing domain then there will be no means for scaling the number of entries needed within routeing databases for all of the addresses of the domain from which such addresses are assigned from. This leads towards the addition of at least one separate AFI for group addresses. Additions to ISO 8348 Addendum 2 provide the international standardization needed to accomplish this objective. In addition in order to minimize the impacts to the Connectionless Network Protocol, the added multicast capabilities are enabled by a Network entity recognizing that a PDU's destination address contains a group NSAP address, thus having such addresses easily recognizable as multicast has implementation benefits.

Research is still on-going in this area and while the use of a separate AFI to accomplish group addressing assignments is recommended, the capabilities are also incorporated (into the additions to ISO 8348 Addendum 2) to permit addressing authorities to allocate multicast addresses within their own domains.

3 Extensions required to the Connectionless Network protocol (ISO 8473).

Annex A provides the detailed changes to ISO 8473 to permit multicast transfer. Only minor changes are required to ISO 8473 to allow multicast PDU transfer. There appears to be no need to propose a new version of this standard since the changes proposed only add optional functionality. The optional functionality is enabled when an Intermediate System with this capability recognizes that a PDU it has received has a group NSAP address in its destination address field. The functionality provided is very similar to that of RFC 1112. The only differences are: (1)The OSI Connectionless Multicast permits Error Reports to be generated due to multicast transfers (with a strong warning provided) while RFC 1112 does not; (2)The OSI Connectionless Multicast permits Source Routeing of PDUs while RFC 1112 does not and (3)The OSI Connectionless Multicast requires that a Network entity originating a PDU send it on every subnetwork to which it is attached while RFC 1112 requires that it be sent to only one. The issues still being considered in utilizing ISO 8473 for multicast transfer are:

a) If "Source Routeing" is used it is assumed that the "multicasting" of the PDU will not take place until the last IS is reached on the list. Further, RFC 1112 does not permit source routes, should this approach be followed?

b) Presently a Network entity sourcing a multicast PDU is required to send the PDU on all subnetworks that the Network entity is attached to. Further work on multicast routeing algorithms is needed to determine if an alternate approach should be used.

4 Extensions required to the Network Service Definition, Addendum 2: Network Layer Addressing (ISO 8348 Addendum 2).

Annex B provides the detailed changes to ISO 8348 Addendum 2 to support multicast transfer. Only minor changes are required to ISO 8348 Addendum 2 to consider the possibility of multiple destination entities. No changes incompatible with the present version of this standard are proposed.

While the detailed changes support the allocation of group NSAP addresses and identify two new AFIs for multicast addresses, considerable flexibility remains for addressing authorities using other AFIs to allocate group addresses. The optional extensions to CLNP are enabled by the recognition of a group NSAP address in the destination address field of a PDU received and thus there are practical reasons to ensure that such addresses are readily detectable.

The two new AFI formats are proposed at this time to enable a very quick determination for a multicast transfer. Annex E contains a further proposal (for use by OSI registration authorities) to embed 48 bit IEEE 802 multicast subnetwork address inside a group NSAP address for use with the AFI's proposed here.

Attention is being paid to the ISO DCC and the ISO 6523-ICD type formats. These appear to have the most potential for use with connectionless multicast. The AFI's proposed here utilize identical syntax as the current AFI's; however, only binary DSP formats are specified since there is no known connectionless multicast transfer requirement where decimal DSP syntax is required or desirable. As work progresses on connection-oriented network techniques (such as X.PMS) further work may be needed on ISO 8348 Addendum 2.

The changes called out to the ISO 8348 Addendum 2 standard are the first step in providing a practical means for administrating group addresses in the OSI domain. The detailed specification of group NSAP address structure needs to be worked by the applicable standard bodies and organizational authorities. It is important that the development of detailed addressing structure for group NSAP addresses be coordinated as is the case for current NSAP address allocations where a common structure was developed by GOSIP, ANSI and the Internet. Specifics on the ISO 8348 Addendum 2 technique covered here are:

a) A Network Entity Group Title (NEGT) identifies a group of Network entities.

b) It is intended that multicast scope be controlled through the structure of the group NSAP address (in combination with the CLNP PDU Lifetime field).

c) All existing rules of NSAP prefix matching hold for a multicast NSAP.

As far as compatibility with the methods discussed in RFC 1112, no major incompatibilities exist. The discussion in RFC 1112 on permanent or transient host groups can be provided in OSI via concepts already in existence within the NSAP address, registration authority and directory concepts and thus this issue is not covered. The automatic translation capability of RFC 1112 for placing the low-order 23-bits of the IP address into the lower order bits of a 48 bit Ethernet multicast address is not provided. As previously mentioned, the Annex E provides a candidate scheme to perform this function utilizing the full 48 bit IEEE 802 multicast subnetwork address.

5 Extensions required to the End system to Intermediate system routeing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service (ISO 8473) - [ISO 9542]

Annex C provides the detailed changes to ISO 9542 to permit multicast transfer. Additional functionality in ISO 9542 is required to ensure that End Systems needing to source and/or sink multicast PDUs can find Intermediate Systems capable of transferring such PDUs and for Intermediate Systems to find which multicast PDUs are needed on which subnetworks. There appears to be no need to propose a new version of this standard since the changes proposed only add optional functionality. The optional functionality is enabled by End Systems passing End System Group Hello (ESGH) PDUs stating intentions of receiving multicast PDUs identified by specific group NSAP addresses contained within the ESGH PDU and by Intermediate Systems passing Intermediate System Group Hello (ISGH) PDUs stating intentions of supporting multicast transfers. The functionality provided is similar to that of RFC 1112 but does use different techniques. A major difference is that the Intermediate Systems in the OSI Connectionless Multicast support a Redirect Function for End Systems which do not know the group SNPA address for a particular group NSAP address, while in RFC 1112 the group SNPA address is directly derived from the IP address.

Note that four methods of providing the optional functionality were considered before the new PDU types (ESGH and ISGH) were selected:

a) No changes to ISO 9542, assuming that all needed differences can be handled by future IS-IS multicast routeing protocol standards and that each IS will determine that Multicast functionality can be determined solely by recognizing group NSAP addresses.

b) Add a second Network Layer Protocol Identifier to PDUs exchanged to support multicast ES-IS functions.

c) Utilize a new Version/Protocol Id Extension identifier to PDUs exchanged to support multicast ES-IS functions.

d) Utilize new PDUs (i.e. ESGH and ISGH) to support multicast ES-IS functions.

Method (a) was not selected since a number of differences were noted between what is done for non-group ES-IS purposes and those needed for group purposes (as listed below). A more conservative approach was taken of assuring that ISs without the multicast capabilities would ignore PDUs using group NSAPs as their destination addresses. By making one of the protocol changes (called out in (b), (c) or (d)) assurance is provided that ISs without multicast capability will ignore PDUs used to support multicast transfer.

Method (b) was not selected since a new Network Layer Protocol Identifier would make the multicast extensions a totally different protocol. The extensions needed are within the scope of the present ISO 9542, these extensions support the routeing of a different type of transfer operation.

Method (c) was not selected since there was no absolute need for a new Version number since all extensions can be provided without impacting existing implementations.

Method (d) was selected since it provided all the needed capabilities with the minimum impact to the standard.

In developing the detailed changes to ISO 9542 a number of differences were noted in the operation of ES-IS operations to support multicast transfer and that which are provided now. In addition to the assumption that multicast routing will be treated separately by ISs when supported, the following differences were noted:

a) Query Configuration (and thus the corresponding Configuration Response function) cannot be requested by ESs. The operation if performed would result in an ES receiving the SNPA address of another ES and not the group SNPA address it is seeking. A system which were to supply such information is actually performing an IS function and the Request Redirect function is available for this purpose.

b) The Request Redirect function must be expanded to support passing the group SNPA address in addition to the other information it might send. An End System receiving such a Redirect must not assume that this is the only station that it is sending to.

c) There is no provision for supporting the Refresh Redirect function for multicast redirects (this is not the primary means assumed of determining group SNPA addresses and thus no new capability is offered at this time).

d) ISs are not permitted to drop Record Configuration operations for group NSAPs used by ESs on subnetworks the IS directly supports. This is due to an assumption on what will be needed by future IS-IS protocols supporting multicast to operate correctly.

e) ESs will report group NSAP addresses only on subnetworks for which it wants to receive such multicast PDUs and not necessarily all subnetworks to which it has SNPAs. Reports on multiple SNPAs is a request for duplicate PDUs.

The issues still being considered in utilizing ISO 9542 for multicast transfer are:

a) Should support be provided for an End System to dynamically find a group SNPA address given knowledge of a group NSAP address? If yes, then a means of implementing the Record Refresh function should be worked out. If no, then the role of the ISs needs to be reconsidered, since the only purpose of the ISGH transfer would be to suggest values for the End System's Configuration Timer.

b) Presently no support is offered to an End System on a subnetwork without an Intermediate System to send a multicast PDU when the group SNPA address is unknown. With no further changes it is implicit that such an End System would drop such PDUs. An option to consider is to have such an End System send to the "All End System Network Entities" SNPA address.

6 Extensions required to the Protocol for providing the connectionless-mode transport service (ISO 8602)

Annex D provides the detailed changes to ISO 8602 to permit multicast transfer. Only very minor changes are required to ISO 8602 to allow multicast PDU transfer. There is no need to propose a new version of this standard since the additional text provided does not add functionality at this layer. The additional capability of multicast transfer is provided by the optional services of the Network layer. While this is out of scope of RFC 1112, applications using RFC 1112 utilize UDP capabilities which are very similar to the technique provided here.

DRAFT

Annex A
Detailed Changes Proposed to CLNP (ISO 8473)

A.1 Change point (a) in 1 (Scope and field of application) to:

"a) procedures for the connectionless transmission of data and control information from one network-entity to one or more peer network- entities."

A.2 Add to the first paragraph, last sentence of 6.1, after the word "destination":

"or destinations"

A.3 Add to the third paragraph, first sentence of 6.1, after the words "particular destination NS User":

"or Users"

A.4 Add to the first paragraph, last sentence of 6.3, right after the first word "If":

"multicast transfer is not supported and if"

A.5 Add a new paragraph 6.3.1:

"6.3.1 Multicast Transfer

The Header Format Analysis function optionally provides capabilities to network-entities which support multicast transfer to supply applicable PDUs directly to end systems served by such a network-entity as well as to forward such PDUs on to other network-entities. This optional functionality is realized through a network-entity with multicast capability identifying a PDU as using multicast transfer via the PDU's NS-Destination-Address field."

A.6 Add to the first paragraph of 6.5 (in three places) after the words "network-entity":

"or network-entities"

A.7 Add a new paragraph at the end of 6.5:

"The Route PDU function optionally provides capabilities to network-entities which support multicast transfer for determining multiple network-entities to which a single PDU shall be forwarded to. This may result in multiple invocations of the Forward PDU function. For PDUs that are received from a different network-entity, the optional functionality for the Route PDU function is realized as a result of the Header Format Analysis function's recognition of the PDU as being a multicast PDU. A network-entity originating a multicast PDU is required to invoke the Forward PDU function for every subnetwork to which it is attached."

A.8 Change the text in the first paragraph of 6.6 from "to identify the "next" system within the subnetwork-specific addressing domain (this may be an intermediate system or the destination end-system" to:

"to identify the "next" system or systems within the subnetwork-specific addressing domain (this may be one or more intermediate system and/or one or more destination end-systems)"

A.9 Add a third note to 6.10.1, at the very end of this section:

"3)The use of the Error Reporting capability must be very carefully controlled in the case of multicast transfers. A multicast PDU with Error Reporting permitted may result in flooding the source network-entity (as well as the networks used) with Error Report PDUs."

NOTE: Annex A of ISO 8473 (Formal Description) was not reviewed.

DRAFT

Annex B

Detailed Changes Proposed to ISO 8348 Addendum 2

B.1 Add a new paragraph 3.4.11:

"3.4.11 group network address: An address which identifies zero or more Network entities which receive NPDUs with this address as the destination."

B.2 Add to the third paragraph, first sentence of 6.1.1 right after "another real subnetwork,":

"a particular group of real end systems on this real subnetwork,"

B.3 Add a new paragraph at the end of 6.1.2:

"In addition to an NSAP address being the information that the Network service provider needs to identify a particular single Network service access point, a group NSAP address identifies a group of access points for different Network entities. The values of the called address of the N-CONNECT primitive and the destination address parameter of the N-UNITDATA primitive are permitted to be group NSAP addresses."

B.4 Table 1:

-Update Table 1 to show two new AFI's assigned from the proper allocation.

B.5 Table 2

-Update Table 2 to show that two new Binary DSP syntax AFI values (??&??) for the ISO DCC Group IDI and ISO 6523-ICD Group IDI formats have been added.

B.6 Add two paragraphs after 8.2.1.2.7:

"8.2.1.2.8 ISO DCC Group IDI format

The IDI is reserved for group Network addresses. The IDI consists of a three-digit numeric code allocated according to ISO 3166. For countries with an ISO member body, the code is assigned to the ISO member body in the country identified by the code. For countries with no ISO member body, the code is assigned to an appropriately sponsored organization in the country identified by the code. The DSP is allocated and assigned by the ISO member body or sponsored organization to which the ISO DCC value has been assigned, or by an organization designated by the holder of the ISO DCC value to carry out this responsibility.

IDP length: 5 digits.

8.2.1.2.9 ISO 6523-ICD Group IDI format

The IDI is reserved for group Network addresses. The IDI consists of a 4-digit International Code Designator (ICD) allocated according to ISO 6523. The ICD identifies an organizational authority responsible for allocating and assigning values of the DSP. The "structure of the code" required by ISO 6523 shall be registered as "According to ISO 8348 Addendum 2".

IDP length: 6 digits."

B.7 Table 3: -Update Table 3 to show that two new Binary DSP syntax AFI values:

“ISO DCC Group..........| __ | 17 | __ | __ |

ISO 6523-ICD Group ...| __ | 17 | __ | __ |”

B.8 Table 5: -Update Table 5 to show that two new Binary DSP syntax AFI values:

“ISO DCC Group ....| Binary | 17 | __ |

ISO 6523-ICD Group ...| Binary | 17 | __ | ”

B.9 Add a new paragraph to the end of Annex A

“A Network Entities Group Title unambiguously identifies a group of Network entities at a particular point in time. Group Network addresses are used to identify either such a title or a group of NSAPs. ”

DRAFT

Annex C

Detailed Changes Proposed to ISO 9542

C.1 Add two sentences to the end of the second paragraph of 5.1 (paragraph beginning with “Configuration information permits...”):

“This information also allows ESs and ISs attached to the same subnetwork and with the optional capability of supporting multicast transfer to dynamically discover each other’s existence and availability. For multicast transfers, configuration information supports ISs with this capability to find out what multicast PDUs are needed by ESs and informs ESs of ISs capable of transferring multicast PDUs that the ES may source. ”

C.2 Add a sentence between the second and last sentences of the last paragraph of 5.1:

“For multicast transfers on subnetworks which provide direct multicast capabilities, route redirection allows ISs to inform ESs of SNPA information which will allow ESs (and ISs) on the same subnetwork to receive multicast PDUs without retransmission. ”

C.3 Add a third bullet to 5.3.1:

“*group SNPA address corresponding to a group NSAP ”

C.4 Change the phrase in the second sentence in the (c) paragraph of 5.4.2.1 “particular NSAP” to:

“particular non-group NSAP ”

C.5 Add a paragraph d) and e) and a Note to 5.4.2.1:

“d) End Systems with the optional OSI multicast capabilities are informed of the reachability and Network Entity Title of each active Intermediate Systems on the subnetwork which have the corresponding OSI multicast capabilities.

e) Intermediate Systems with the optional OSI multicast capabilities are informed of the group NSAP addresses which End Systems on this subnetwork need to receive.

Note: Intermediate Systems with the optional OSI multicast capabilities do receive information identifying which particular ESs on the broadcast network want PDUs with particular group NSAP addresses as their destination address; however, the critical information is which multicast PDUs are needed not which ESs need them. Future routeing standards may (or may not) make use of the fact that the ISs can identify which ESs need particular multicast PDUs. ”

C.6 Add sections 6.2.3 and 6.2.4:

"6.2.3 Report Group Configurations by End Systems

An End System which needs to receive or continue to receive any multicast PDUs (i.e. NPDUs with group NSAP addresses as their destination address), constructs and transmits ESGH PDUs to inform Intermediate Systems with the optional multicast capabilities of the multicast PDUs it needs to receive. This may be done by constructing ESGH PDUs for each group NSAP address. Alternatively, ESGH PDUs may be constructed which convey information about more than one NSAP address at a time, up to the limits imposed by the permitted SNSDU size and the maximum header size of the ESGH PDU. Each ESGH PDU is transmitted by issuing an SN-UNITDATA.Request with the following parameters:

SN_Userdata (SNSDU) <- ESGH PDU

SN_Destination _Address <- multi-destination address that indicates "All Intermediate System Network Entities"

Where an End System supports more than one SNPA the information about each group NSAP address desired for receiving on a particular subnetwork serving the End System shall be transmitted on that subnet via its SNPA. It is permissible for an End System to report group NSAPs on multiple SNPAs; however, duplicate multicast PDUs should be anticipated.

Timer considerations are identical with the Report Configuration by End System function. Timing considerations for ESHs are identical to those for ESGHs.

6.2.4 Report Group Configuration by Intermediate Systems

An Intermediate System with the optional multicast capabilities constructs a single ISGH PDU containing the IS's Network entity title and issues one SN_UNITDATA.Request on each SNPA to which it is attached with the following parameters:

SN_Userdata (SNSDU) <- ISGH PDU

SN_Destination_Address <- multi-destination address that indicates "All End System Network Entities"

Timer considerations are identical with the Report Configuration by Intermediate System function. Timing considerations for ISHs are identical to those for ISGHs. "

C.7 Section 6.3, Change the two occurrences of "ESH or ISH" to:

"ESH, ESGH, ISH or ISGH"

C.8 Section 6.3.1, Add a second paragraph:

"On receipt of an ESGH PDU an IS with the optional multicast capabilities extracts the configuration information and stores the {group NSAP address, SNPA} in its local multicast routeing information base replacing any other information for the same entry. "

C.9 Add a paragraph after the first paragraph in C.3.2:

"On receipt of an ISGH PDU an ES with the optional multicast capabilities extracts the configuration information and stores the {NET, SNPA} into its routeing information base, flagging it as an IS capable of transferring multicast PDUs and replacing any other information for the same entry. If insufficient space is available to store the new configuration information the PDU is discarded.

C.10 Change “ISH” in the present second paragraph of 6.3.2 to:

“ISH or ISGH ”

C.11 Add to paragraph (c) of 6.5 prior to the word “NSAP”:

“non-group ”

C.12 Add a Note under point (c) of 6.5:

“Note: The Query Configuration function cannot be performed to find the corresponding SNPA address of a group NSAP address since the addressing information needed is the corresponding group SNPA address and not the SNPA address of a particular End System responding. On a large broadcast subnetwork, many different Configuration Responses could result each incorporating a different End System Address. For a system to supply the multicast SNPA address needed, a system would actually be performing an Intermediate System function and in this case the Request Redirect function is available. ”

C.13 Add a paragraph prior to the last paragraph of 6.7:

“A system with the optional multicast capabilities which chooses to implement this function executes it on detecting, by receiving an ESGH or ISGH PDU, that another system has just become available. It then constructs an ISGH or ESGH PDU respectively as described in 6.2.4 or 6.2.3, but transmits it specifically addressed to the newly operational system using an SN_UNITDATA.Request with the following parameters:

SN_Userdata <- ESGH or ISGH PDU

SN_Destination_Address <- SN_Source_Address parameter value from the Sn_UNITDATA.Indication containing the original ESGH or ISGH PDU as its SN_Userdata. ”

C.14 Add to the end of point (c) of 6.8 along with a new point (d):

c)..., or

d) In the case of a group NSAP, the group subnetwork address. ”

C.15 Note that Redirect information will eventually time-out and that the present Refresh Redirect Function will not maintain it. This is not the normal means anticipated for determining group SNPA addresses by End Systems and no means of maintaining this is provided at this time.

C.16 Add two new entries to Table 2, determine actual binary values at a later time:

```
“       ESGH PDU a b c d e
        ISGH PDU f g h i j”
```

C.17 Throughout section 7.3, including the subparagraph and figures and section titles, replace:

“ESH” with “ESH and ESGH” and “ISH” with “ISH and ISGH”; {detailed changes may require the word “and” to be replaced by a “,” or the word “or”}

C.18 Add a last sentence to the end of the first paragraph of 7.3.6:

"In the case of a multicast transfer, the Subnetwork Address Parameter is used to indicate the group subnetwork address corresponding to the group NSAP address that an End System placed in the Destination Address parameter sent directly to the Intermediate System."

C.19 Add sections 7.8 and 7.9:

"7.8 End System Group Hello (ESGH) PDU

The ESGH PDU has the format shown in figure 9.

7.9 Intermediate System Group Hello (ISGH) PDU

The ISGH PDU has the format shown in figure 10."

C.20 Change the title of Figure 9 to:

"ESH and ESGH PDU Format"

C.21 Change the title of Figure 10 to:

"ISH and ISGH PDU Format"

C.22 Change the title of Figure 12 to:

"RD PDU Format when Redirect to an ES or to a group subnetwork address"

C.23 Changes to the Conformance (section 8) are needed. Changes to the Static Conformance Requirements (8.1, 8.1.1, 8.1.2) and the Protocol Implementation Conformance Statement will be supplied at a later time. All capabilities to support multicast transfer are to be optional.

DRAFT

Annex D

Detailed Changes Proposed to ISO 8602

D.1 Change the end of the first point (a) under paragraph 1 from “to one peer transport entity;” to:

“to one or more peer transport entities;”

D.2 Add to the end of 3.3.2:

“The destination-transport address may identify a group of Transport Service users connected to different network entities depending on the services used and provided by the network service provider.”

D.3 Add a new sentence at the end of 5.2:

“Depending on the services provided by the Network service, a transport user may be able to send data to a group of other transport users and receive PDUs intended for a group of transport users via the use of the Destination address parameters in Table 3.”

DRAFT

Annex E

Proposed American National Standard

Structure and Semantics of the Domain Specific Part (DSP) of the OSI Group Network Service Access Point Address

1. Scope, Purpose and Field of Application

1.1 Scope. The standard for OSI Network layer addressing, ISO 8348 Addendum 2, does not specify the structure or the semantics of the Domain Specific Part of the NSAP address; this is left to the individual Registration Authorities which are responsible for allocation and assignment of NSAP addresses in the formats identified by individual values of the Authority and Format Identifier (AFI) and Initial Domain Identifier (IDI).

The ISO community has started work in developing the standards (and modifications to existing standards) to support the group addresses for multicast transfer of PDUs. One of the formats under consideration is the ISO DCC Group AFI. In the case of the AFI/IDI value??/840 (AFI = ??), IDI = 840), the responsible Registration Authority is ANSI. This American National Standard specifies the structure and semantics of the DSP part of the NSAP address when the value of the AFI and IDI parts of the address are?? and 840 (respectively).

1.2 Purpose.

The OSI Network layer routing architecture, as specified by ISO Technical Report 9575, defines a hierarchy of routing areas and domains. The OSI routing protocols (ISO 9542, ISO/IEC 10030, ISO/IEC 10589, and the working draft inter-domain routing protocol) have been specified in such a way that routing operations are much more efficient if the hierarchical structure of routing areas and domains is reflected by a corresponding hierarchical structure of the OSI NSAP address, including a hierarchical sub-structure of the DSP. Such a structure is strongly implied by, in particular, ISO 10589, and has been adopted by the United States Government OSI Profile (GOSIP) specification. The purpose of this [proposed] American National Standard is to specify the structure and semantics of the DSP for those group NSAP addresses over which ANSI has direct authority. In addition, there are many advantages obtainable if other addressing authorities structure group NSAPs in a consistent way; however, while this proposed standard is intended to support discussions in the development of ISO group addressing concepts and standards, this proposed standard is intended only for ANSI.

1.3 Field of Application. This American National Standard applies to the allocation and assignment of OSI NSAP addresses when the value of the AFI and IDI parts of the address are?? and 840 (respectively). The field of application includes all environments in which this class of NSAP addresses is used, whether or not the hierarchical routing architecture to which clause 1.2 refers is used.

The DSP structure defined by this standard may also be used in other circumstances that are not specifically included within the scope and field application of this standard.

3 Definitions (incorporate section 3 of X3S3.3/91-151 “Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address”)

4 Abbreviations (incorporate section 4 of X3S3.3/91-151 “Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address”)

5 Overview of the NSAP Address (incorporate section 5 of X3S3.3/91-151 “Structure and Semantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address”)

6 Structure and Semantics of the DSP (incorporate section 6 of X3S3.3/91-151 “Structure and Se-

mantics of the Domain Specific Part (DSP) of the OSI Network Service Access Point (NSAP) Address") with the following changes and additions:

{different Figure 1}

IDP		DSP						
AFI	IDI	DFI						
??	840	220	org	res	rd	area	IEEE 802 SNPA address	sel
1	2	1	3	2	2	2	6	1

Figure 1 -- Structure of the DSP

{CHANGE 6.2.2}

6.2.2 [DFI] The DSP Format Identifier, which specifies the version of this DSP structure Standard. The value of this field shall be binary 1101 1010 (shown in Figure 1 as decimal 220).

{Change 6.2.7}

6.2.7 [system] This field shall utilize a valid IEEE 802 SNPA address administered by the IEEE (i.e. the Universal/Local bit set to Universal) and the Group addressing selected (i.e. the Individual/Group bit set to Group). The IEEE 802.2 (ISO 8802.2) LSAP address shall not be incorporated into this field.

{Add a Note at the end of the present text (after 6.2.8)}

Note: a group NSAP under the scope and field of application of this standard with org=0 and res=0 and rd=0 and area=0 and sel=0 has a one to one correspondence with the IEEE 802 SNPA Address used in the system field and provides no additional addressing information beyond this SNPA address.

{Add the following paragraphs}:

6.3 SNPA Requirements. A network entity utilizing a subnetwork which supports IEEE 802 forty eight bit SNPA addressing shall use the contents of the DSP system field as the subnetwork destination address used for PDU transfer under the scope and field of application of this standard.