

Host Group Extensions for CLNP Multicasting

Status of this Memo

This memo defines an Experimental Protocol for the Internet community. This memo does not specify an Internet standard of any kind. Discussion and suggestions for improvement are requested. Distribution of this memo is unlimited.

Abstract

This memo documents work performed in the TUBA (TCP/UDP over Bigger Addresses) working group of IPng area prior to the July 1994 decision to utilize SIPP-16 as the basis for IPng. The TUBA group worked on extending the Internet Protocol suite by the use of ISO 8473 (CLNP) and its related routing protocols. This memo describes multicast extensions to CLNP and its related routing protocols for Internet multicast use. Publication of this memo does not imply acceptance by any IETF Working Group for the ideas expressed within.

This memo provides a specification for multicast extensions to the CLNP protocol similar to those provided to IP by RFC1112. These extensions are intended to provide the mechanisms needed by a host for multicasting in a CLNP based Internet. This memo covers addressing extensions to the CLNP addressing structure, extensions to the CLNP protocol and extensions to the ES-IS protocol. An appendix discusses the differences between IP multicast and the CLNP multicast approach provided in this memo.

Acknowledgments

The specification provided here was developed by a number of individuals in the IETF TUBA working group as well as the ANSI X3S3.3 and ISO SC6 WG2 committees. Key contributions were made by Steve Deering, Joel Halpern, Dave Katz and Dave Oran.

Table of Contents

| | |
|--|----|
| 1. Introduction | 2 |
| 2. Levels of Conformance..... | 3 |
| 3. Group Network Addresses..... | 4 |
| 4. Model of a CLNP End System Multicast Implementation.... | 8 |
| 5. Extensions to the CLNP Protocol..... | 8 |
| 6. Extensions to the ES-IS Routeing Protocol | 15 |
| 7. Security Considerations | 39 |
| Appendix A. Differences with RFC 1112 | 40 |
| Appendix B. Issues Under Study | 43 |
| References | 44 |
| Author's Address | 45 |

1. Introduction

This memo provides a specification for multicast extensions for CLNP in order to provide a CLNP based Internet the capabilities provided for IP by RFC 1112 (Host Extensions for IP Multicasting) [RFC1112]. This memo uses an outline similar to that of RFC 1112.

Paraphrasing RFC 1112, "CLNP multicasting is the transmission of a CLNP datagram to a "host group", a set of zero or more End Systems identified by a single group Network address (GNA). A multicast datagram is delivered to all members of its destination host group with the same "best-efforts" reliability as regular unicast CLNP datagrams, i.e., the datagram is not guaranteed to arrive intact at all members of the destination group or in the same order relative to other datagrams.

"The membership of a host group is dynamic; that is End Systems may join and leave groups at any time. There is no restrictions on the location or number of members in a host group. An End System may be a member of more than one group at a time. An End System need not be a member of a group to send datagrams to it.

"A host group may be permanent or transient. A permanent group has an administratively assigned GNA. It is the address, not the membership of the group, that is permanent; at any time a permanent group may have any number of members, even zero.

"Internetwork forwarding of CLNP multicast datagrams is handled by "multicast capable" Intermediate Systems which may be co-resident with unicast capable Intermediate Systems.

The multicast extensions to the CLNP addressing structure defines group Network addresses which identify host groups. The multicast extensions to CLNP provides a means for identifying a CLNP packet and

provides scope control mechanisms for CLNP multicast packets. The multicast extensions to the ES-IS protocol provide the mechanisms needed for a host to exchange control information with multicast capable routers. These extensions to the ES-IS protocol provide for a host to "announce" which multicast packets are of interest and for a multicast capable router to dynamically "map" group Network addresses to subnetwork addresses.

This memo specifies the extensions required by an End System to make use of CLNP multicast. In addition the requirements placed upon multicast capable Intermediate Systems to exchange information with multicast capable End Systems is specified. No specifications are provided related to the information exchanges between Intermediate Systems to support multicast route selection or multicast Protocol Data Unit (PDU) forwarding. A discussion of multicast route selection and PDU forwarding has been written by Steve Deering [Deering91]. Note that for these multicast extensions to work there must exist an uninterrupted path of multicast capable routers between the End Systems comprising a host group (such paths may utilize tunneling (i.e., unicast CLNP encapsulated paths between multicast capable CLNP routers)). In order to support multicast route selection and forwarding for a CLNP based internet additional specifications are needed. Specifications of this type could come in the form of new protocols, extensions to the current CLNP based routing protocols or use of a technique out of the IETF's Inter-Domain Multicast Routing (IDMR) group. The IDMR group is currently investigating multicast protocols for routers which utilize a router's unicast routing protocols, this approach may extend directly to CLNP routers.

While many of the techniques and assumptions of IP multicasting (as discussed in RFC 1112) are used in CLNP multicasting, there are number of differences. Appendix A describes the differences between CLNP multicasting and IP multicasting. This memo describes techniques brought in directly from projects within ISO to incorporate multicast transmission capabilities into CLNP [MULT-AMDS].

2. Levels of Conformance

There are three levels of conformance for End Systems to this specification:

Level 0: no support for CLNP multicasting.

There is no requirement for a CLNP End System (or Intermediate System) to support CLNP multicasting. Level 0 hosts should be unaffected by the presence of multicast activity. The destination addresses used in support of multicast transfers, the GNA, should not be enabled by a non-multicast capable End System and the PDUs

themselves are marked differently than unicast PDUs and thus should be quietly discarded.

Level 1: support for sending but not receiving CLNP multicast PDUs.

An End System originating multicast PDUs is required to know whether a multicast capable Intermediate System is attached to the subnetwork(s) that it originates multicast PDUs (i.e., to determine the destination SNPA (subnet) address). An End System with Level 1 conformance is required to implement all parts of this specification except for those supporting only Multicast Announcement. An End System is not required to know the current Multicast Address Mapping to start originating multicast PDUs.

Note: It is possible for End System not implementing Multicast Address Mapping to successfully originate multicast PDUs (but with the End System knowing of the existence of a multicast capable Intermediate System). Such operation may lead to inefficient subnetworks use. Thus when an End System continues (or may continue) to originate multicast PDUs destined for the same group, implementations are to provide Multicast Address Mapping support.

Level 2: full support for CLNP multicasting.

Level 2 allows a host to join and leave host groups as well as send CLNP PDUs to host groups. It requires implementation by the End System of all parts of this specification.

3. Group Network Addresses

Individual Network addresses used by CLNP for End System addressing are called Network Service Access Points (NSAPs). RFC 1237 defines the NSAP address for use in the Internet. In order to provide an address for a group of End Systems, this specification does not change the definition of the NSAP address, but adds a new type of identifier - the group Network address - that supports a multicast Network service (i.e., between a single source NSAP, identified by an individual Network address, and a group of destination NSAPs, identified by a group Network address). Host groups are identified by group Network addresses.

In the development of multicast address extensions to CLNP, requirements were identified for: (1) "easily distinguishing" group addresses at the Network layer from NSAP addresses; (2) leaving the currently allocated address families unaffected and (3) ensuring that the approach taken would not require the establishment of new addressing authorities. In addition, it was agreed that providing multicast options for all OSI Network layer users was desirable and

thus the group Network addressing solution should support options for all address formats covered by ISO/IEC 8348 | CCITT Recommendation X.213. The only viable means identified for meeting all requirements was via creating a new set of AFI values with a fixed one-to-one mapping between each of the existing AFI values and a corresponding group AFI value.

Group Network addresses are defined by creating a new set of AFI values, one for each existing AFI value, and a fixed one-to-one mapping between each of the existing AFI values and a corresponding group AFI value. The syntax of a group Network address is identical to the syntax of an individual Network address, except that the value of the AFI in an individual Network address may be only one of the values already allocated for individual Network addresses, whereas the value of the AFI in a group Network address may be only one of the values allocated here for group Network addresses. The AFI values allocated for group Network addresses have been chosen in such a way that they do not overlap, in the preferred encoding defined by ISO/IEC 8348 | CCITT Recommendation X.213, with any of the AFI values that have already been allocated for individual Network addresses.

3.1 Definitions

group Network address: an address that identifies a set of zero or more Network service access points; these may belong to multiple Network entities, in different End Systems.

individual Network address: an address that identifies a single NSAP.

3.2 CLNP Addresses

A discussion of the CLNP address format is contained in RFC 1237. The structure of all CLNP addresses is divided into two parts the Initial Domain Part (IDP) and the Domain Specific Part (DSP). The first two octets of the IDP are the Authority and Format Identifier (AFI) field. The AFI has an abstract syntax of two hexadecimal digits with a value in the range of 00 to FF. In addition to identifying the address authority responsible for allocating a particular address and the format of the address, the AFI also identifies whether an address is an individual Network address or a group Network address. There are 90 possible AFI values to support individual Network address allocations. In addition, when the AFI value starts with the value "0" this identifies that the field contains an incomplete individual Network address (i.e., identifies an escape code).

Table 1 allocates 90 possible AFI values to support group Network address allocations. In addition if the first two digits of the IDP are hexadecimal FF, this indicates the presence of an incomplete

group Network address. The allocation of group addresses is restricted to be only from the AFI values allocated for the assignment of group addresses in Table 1. An addressing authority in allocating either Network addresses or authorizing one or more authorities to allocate addresses, allocates both individual and the corresponding group addresses. Thus each block of addresses allocated by an addressing authority (or its sub-authority) contains a block of individual Network addresses and group Network addresses. The individual and group address block allocated are differentiated by the AFI values used which are related as shown in Table 1.

Group Network addresses are only used as the destination address parameter of a CLNP PDU. Source Address parameters are never permitted to be group Network addresses.

Table 2 lists the AFI values which have not been assigned, at this time, for the support of neither individual nor group address allocation. Future assignment of these AFI values is possible. Additional information concerning individual Network addresses (i.e., NSAP and NET (Network Entity Titles)) is contained in RFC 1237.

Note: While the format of the Initial Domain Part of a group Network address is assigned, the format for the Domain Specific Part of the group Network address is specified by an addressing authority and is out of the scope of this memo. While NSAP address assignments are typically made to support hierarchical unicast routing, a similar consideration for group Network address assignments may not exist.

TABLE 1 - Relationship of AFI Individual and Group Values

| Individual | Group | Individual | Group | Individual | Group |
|------------|-------|------------|-------|------------|-------|
| 0x | FF | | | | |
| 10 | A0 | 40 | BE | 70 | DC |
| 11 | A1 | 41 | BF | 71 | DD |
| 12 | A2 | 42 | C0 | 72 | DE |
| 13 | A3 | 43 | C1 | 73 | DF |
| 14 | A4 | 44 | C2 | 74 | E0 |
| 15 | A5 | 45 | C3 | 75 | E1 |
| 16 | A6 | 46 | C4 | 76 | E2 |
| 17 | A7 | 47 | C5 | 77 | E3 |
| 18 | A8 | 48 | C6 | 78 | E4 |
| 19 | A9 | 49 | C7 | 79 | E5 |
| 20 | AA | 50 | C8 | 80 | E6 |
| 21 | AB | 51 | C9 | 81 | E7 |
| 22 | AC | 52 | CA | 82 | E8 |
| 23 | AD | 53 | CB | 83 | E9 |
| 24 | AE | 54 | CC | 84 | EA |
| 25 | AF | 55 | CD | 85 | EB |
| 26 | B0 | 56 | CE | 86 | EC |
| 27 | B1 | 57 | CF | 87 | ED |
| 28 | B2 | 58 | D0 | 88 | EE |
| 29 | B3 | 59 | D1 | 89 | EF |
| 30 | B4 | 60 | D2 | 90 | F0 |
| 31 | B5 | 61 | D3 | 91 | F1 |
| 32 | B6 | 62 | D4 | 92 | F2 |
| 33 | B7 | 63 | D5 | 93 | F3 |
| 34 | B8 | 64 | D6 | 94 | F4 |
| 35 | B9 | 65 | D7 | 95 | F5 |
| 36 | BA | 66 | D8 | 96 | F6 |
| 37 | BB | 67 | D9 | 97 | F7 |
| 38 | BC | 68 | DA | 98 | F8 |
| 39 | BD | 69 | DB | 99 | F9 |

TABLE 2 - AFI values reserved for future allocation

| |
|-------|
| 1A-1F |
| 2A-2F |
| 3A-3F |
| 4A-4F |
| 5A-5F |
| 6A-6F |
| 7A-7F |
| 8A-8F |
| 9A-9F |
| FA-FE |

4. Model of a CLNP End System Multicast Implementation

The use of multicast transmission by a CLNP End System involves extensions to two protocols: CLNP and the ES-IS Routing Protocol. To provide level 0 service (no support for CLNP multicast), no extensions to these two protocols are required. To provide level 1 service (support for sending but not receiving CLNP multicast PDUs) all extensions contained in the following sections are required except for those supporting only Multicast Announcement. In order to support level 2 service (full support for CLNP multicasting), the extensions contained in the following sections are required. Extensions identified for Intermediate Systems are not required (or appropriate) for End Systems. Multicast transmission also requires the use of a group Network address (as previously described) as the destination address parameter.

5. Extensions to the CLNP protocol

This section provides extensions to the CLNP Protocol [CLNP] ISO 8473-1, to support multicast transmission. These additions provide procedures for the connectionless transmission of data and control information from one network-entity to one or more peer network-entities.

In developing the multicast extensions for CLNP a decision was needed on how to "mark" a packet as multicast (versus the current unicast packets). Such marking is necessary since the forwarding behavior for multicast packets is different (e.g., multiple copies of a packet may need to be forwarded). The two alternatives considered were to mark the packet (via a particular field) or to mark the destination address, in the end both were done. The destination address for a multicast PDU identifies a host group which is of a very different nature than the unicast NSAP address. Rather than changing the

nature of NSAP addresses, a new set of addresses were created named group Network addresses which are marked within the first octet (i.e., the AFI field) with values reserved for group Network addresses.

Consideration was given to no further marking of the PDU; however, a problem was identified with only using the group Network address to identify multicast packets. Currently routers implementing the IS-IS Intra-Domain protocol as Level 1 routers when receiving a packet with an unknown destination address are permitted to either discard the packet or send it to a Level 2 router. Such actions by non-multicast capable routers to multicast packets can lead to non-deterministic behavior. Level 1 routers upon receiving a packet containing a group Network address might pass the packet up to a Level 2 router (which may or may not be multicast capable) or it might discard it. Depending upon the circumstances this might lead to whole regions missing packets or packet duplication (possibly even explosion). The result was to seek deterministic behavior by non-multicast capable routers by creating a new PDU type (Multicast Data PDU) and inserting into the CLNP reasons for discard: receiving a PDU of unknown type. Note that this reason for discard is mandatory on multicast capable and non-multicast capable CLNP implementations.

5.1 Definitions

multicast: Data transmission to one or more destinations in a selected group in a single service invocation.

multicast capable Intermediate System: An Intermediate System which incorporates the multicast features of the Network layer.

5.2 Addresses

The destination address parameter of a multicast PDU shall contain a group Network address. The source address parameter shall be an individual Network address.

5.3 Extensions to the current protocol functions

In order to support multicast transmissions the following optional CLNP protocol functions will be implemented:

5.3.1 Header Format Analysis function

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 type and the PDU's destination address field.

If a Network entity supports multicast transmission, then the header format analysis function shall provide checking to ensure that a PDU does not contain a group Network address in the source address field. Any PDU header analyzed to have a group address in the source address field shall be discarded.

5.3.2 Route PDU function

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 and hence the need to make multiple copies of the PDU. 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 attached to more than one subnetwork when originating a multicast PDU is permitted to originate the PDU on more than one subnetwork.

Note: The ES-IS function "Extensions to the ISO CLNP Route Function by End Systems" discussed in section 6.10 identifies on which subnetworks an End System attached to more than one subnetwork must originate multicast PDUs on.

Note: The purpose in allowing an originating Network entity to originate a multicast PDU on multiple subnetworks is to support the development of multicast IS-IS protocols which will need to determine on which subnetworks a multicast PDU has visited. This behavior is predicated on the assumption that the Intermediate Systems in the OSI environment performing multicast forwarding form a connected set.

5.3.3 Forward PDU function

This function issues an SN-UNITDATA request primitive, supplying the subnetwork or Subnetwork Dependent Convergence Function (SND CF) identified by the route PDU function with the protocol data unit as user data to be transmitted, the address information required by that subnetwork or SND CF to identify the "next" system or systems within the subnetwork-specific addressing domain (this may be one or more Intermediate Systems and/or one or more destination End Systems), and quality of service constraints (if any) to be considered in the processing of the user data.

5.3.4 Discard PDU function

Add an additional reason for discard - a PDU is received with an unknown type code.

5.3.5 Error reporting function

It is important to carefully control the use of the error reporting capability in the case of multicast transfers. The primary concern is to avoid the occurrence of broadcast storms and thus a multicast PDU may not cause the origination of another multicast PDU. This is the primary reason that the source address is not permitted to be a group address. In addition, a multicast PDU with error reporting permitted can result in flooding the source network-entity (as well as the networks used) with Error Report PDUs.

While error reports are permitted on multicast PDUs, a PDU with a group Network address in the source address field shall not be responded to with an Error Report. This is to ensure that a multicast PDU does not generate another multicast PDU. If the source address is identified as a group address then an error report PDU shall not be generated and the original PDU shall be discarded.

5.3.6 Source routing functions

No source routing capability is provided for multicast PDU transfer. The NS provider shall not accept a multicast PDU with source route parameters.

5.4 Scope control function

5.4.1 Overview

The scope control function is an option for multicast PDU forwarding only. The scope control function allows the originator to limit the forwarding of the multicast PDU. The scope control function provides the capability to limit the relaying of a particular PDU based on the individual Network addressing hierarchy and/or limit the amount of multicast expansion which can take place. In cases where both forms of scope control are applied to the same PDU, forwarding will cease once either has reached its scope control limit.

5.4.2 Prefix Based Scope Control

The prefix based scope control function allows the originator to specify a specific set of address prefixes where the multicast forwarding of a PDU by an Intermediate System occurs only if one of the prefixes matches the Network Entity Title (NET) of the

Intermediate System. Prefix based scope control may be selected only by the originator of a PDU. Prefix based scope control is accomplished using one or more address prefixes held in a parameter within the options part of the PDU header. The length of this parameter is determined by the originating network entity, and does not change during the lifetime of a PDU.

When an Intermediate System receives a multicast PDU containing a prefix based scope control parameter, forwarding is only performed if every octet of one of the prefixes contained in the prefix based scope control parameter matches that Intermediate System's NET, starting from the beginning of its NET. If no such prefix match exists, the Intermediate System discards the PDU. The error reporting function shall not be invoked upon PDU discard.

5.4.3 Radius Scope Control

The radius scope control function allows the originator to specify a maximum logical distance where multicast expansion can occur. It is closely associated with the header format analysis function. Each IS receiving a multicast PDU which is capable of expanding and which contains a Radius Scope Control parameter, decrements the Radius Scope Control field in the PDU by an administratively set amount between 0 and the maximum value of the field. An IS, when it decrements the Radius Scope Control field, shall place a value of 0 into this field if its current value is less than the amount it is to decrement by. This function determines whether the PDU received may be forwarded or whether its Radius has been reached, in which case it shall be discarded. An Intermediate System shall not forward a multicast PDU containing a Radius Scope Control parameter with a value of 0. The error reporting function shall not be invoked upon PDU discard.

5.4.3.1 Radius Scope Control Example

The Radius Scope Control parameter is useful where policies have been established across the potential forwarding path. One possible policy for Internet use is for multicast capable routers to treat this field as a hop count within a domain (decrement by one unit) and for inter-domain routers to either decrement this field to an even multiple of 256 when crossing domains where prior agreements have been made or decrement this field to 0 (i.e., discard the packet) for other domains.

5.5 Structure and Encoding of PDUs

Multicast transmission is accomplished via the transfer of Multicast Data (MD) PDUs. The PDU type code for a MD PDU is "1 1 1 0 1". The format of the MD PDU is identical to that of the Data (DT) PDU. The MD and DT PDU may contain the same optional parameters with the following exceptions: (1)The source routing parameter is permitted within DT PDUs but not MD PDUs; and (2)The scope control parameter is permitted within MD PDUs but not DT PDUs.

5.6 Optional parameters for multicast support

5.6.1 Prefix Based Scope Control

The prefix based scope control parameter specifies one or more address prefixes for which Intermediate System forwarding requires a match of one of the contained prefixes with the beginning of the Intermediate System's NET.

Parameter Code: 1100 0100

Parameter Length: variable

Parameter Value: a concatenation of address prefix entries

The parameter value contains an address prefix list. The list consists of variable length address prefix entries. The first octet of each entry gives the length of the address prefix denominated in bits that comprises the remainder of the entry. If the length field does not specify an integral number of octets then the prefix entry is followed by enough trailing zeroes to make the end of the entry fall on an octet boundary. The list must contain at least one entry.

The prefix shall end on a boundary that is legal in the abstract syntax of the address family from which it is derived. For example, the encoding of a prefix whose DSP is expressed in decimal syntax must end on a semi-octet boundary, while the encoding of a prefix whose DSP is expressed in binary syntax can end on an arbitrary bit boundary. If the end of the prefix falls within the IDP, then the prefix must end on a semi-octet boundary and must not contain any padding characters.

Note: The length of the prefix based scope control parameter is determined by the originator of the PDU and is not changed during the lifetime of the PDU.

5.6.1.1 Prefix matching

A prefix that extends into the DSP shall be compared directly against the encoded NET address, including any padding characters that may be present. A prefix which does not extend into the DSP shall be compared against the derived quantity NET', which is obtained from the NET address by removing all padding characters (as defined by the binary encoding process of ISO 8348).

The existence of a match shall be determined as follows:

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

5.6.2 Radius Scope Control

The radius scope control parameter specifies the logical distance that a multicast PDU can be forwarded.

Parameter Code: 1100 0110

Parameter Length: two octets

Parameter Value: two octets which represents the remaining distance, that the PDU can be forwarded, in administratively set units.

5.7 Provision of the Underlying Service

For a subnetwork that provides an inherent multicast capability, it is the functionality of the SND CF to provide the mapping between group Network addresses and the corresponding addressing capability of the subnetwork.

5.8 Conformance

All of the extensions provided to the functions to support multicast capability are optional. For an End System or Intermediate System which is not multicast capable these extensions are not applicable. An implementation claiming conformance as a multicast capable End System shall meet all of the requirements for an End System which is not multicast capable and also provide all of the multicast extensions provided here. An implementation claiming conformance as a

multicast capable Intermediate System shall meet all of the requirements for an Intermediate System which is not multicast capable and also provide all of the multicast extensions provided here.

6. Extensions to the ES-IS Routeing Protocol

This section provides optional extensions to the ES-IS Routeing Protocol [ES-IS], ISO 9542 to support the transfer of multicast PDUs. It is an explicit goal of this specification that ESs and ISSs, some of which will have multicast capabilities and some without, will be able to fully function on the same subnetworks. This specification does not change any aspect of a currently defined (i.e., non-multicast) ISO 9542 implementation, it adds new optional functionality not modifying current functionality. Two basic functions are provided: multicast announcement and multicast address mapping.

6.1 Overview of the protocol

6.1.1 Operation of ESs receiving multicast PDUs

ESs, upon initialization and periodically thereafter, will construct End System Group Hello (ESGH) PDUs identifying, by particular group Network addresses, the multicast PDUs it wishes to receive. The ES will periodically originate (announce) these ESGH PDUs on the subnetwork it wishes to receive these on. Reporting the same group Network address on multiple subnetworks may result in the reception of duplicate PDUs. ES-IS operations related to requesting the same group Network address on multiple subnetworks are handled totally independently (e.g., using different logical timers,...). It is permitted for an ES to report a number of group Network addresses in the same ESGH PDU. The only restrictions placed on providing multiple group Network addresses within the same ESGH PDU are that all packets requested are to be received on the same subnet, with the same holding time and that the ESGH PDU be of length equal to or less than its maximum packet size constraint. Note that each group Network address in the ESGH PDU is paired with its own SNPA (subnetwork point of attachment) address.

An ES will always have an SNPA address associated with each of its active group Network addresses. An SNPA address is a subnetwork address, in the case of a subnetwork which uses IEEE 802 addresses the SNPA address is a 48 bit IEEE 802 MAC (media access control) address. Of particular interest is the address used to mark the destination group. For a subnetwork using IEEE 802 addressing a group SNPA address uses a particular bit position to "mark" group SNPA addresses.

Upon initialization the ES may have static SNPA address associations (Pre-configured SNPA addresses). For any group Network address without a Pre-configured SNPA address that the ES wishes to receive, the ES will associate the "All Multicast Capable End Systems" SNPA address. Upon receiving a Multicast Address Mapping (MAM) PDU containing a group Network address that the ES is announcing, the ES will use the SNPA address pairing contained in the MAM PDU for that group Network address. Upon the expiration of the Mapping Holding Timer, the ES shall revert back to associating either the Pre-configured SNPA address if one exists or the "All Multicast Capable End Systems" SNPA address for the specific group Network address. While an ES is permitted to listen in on other ESs announcements (needed for the damping option), an ES is not permitted to change its group Network address to SNPA address mapping based on the announcement of other ESs.

Optionally, the ES may perform damping (resetting a Multicast Announcement Timer corresponding to a particular group Network address) if the conditions necessary to withhold a particular announcement are met. In order to perform damping the following conditions must be met: (1) The ES must be processing other ES's announcements; (2) An ESGH PDU is received that identifies the exact same group Network address and SNPA address pairing on a particular subnetwork that this ES is announcing on; (3) The Multicast Holding Timer parameter value in the ESGH PDU received is equal to or greater than the Multicast Holding Timer value, for this subnetwork, that is being used by the ES processing this ESGH PDU.

ESs will utilize a local default value for their Multicast Announcement Timer to control the period for sending out their ESGH PDUs. The Active Multicast IS, if one exists on a particular subnetwork, may suggest a value for ESs on the subnetwork to use for their Multicast Announcement Timer for a specific group Network address. In order to support the optional damping function, ESs are required to incorporate a 25% jittering to the timer values that they are using.

6.1.2 Operation of ESs originating multicast PDUs

The ES originating multicast packets identified by a specific group Network address is not required to be a receiver of such packets (and thus is not announcing that particular group Network address). The origination of multicast PDUs involves two differences to the origination of unicast PDUs. The two differences are: (1) The mechanism for selecting a destination SNPA address and (2) For End Systems attached to more than one subnet, the decision on which subnet(s) to originate the PDUs.

The destination SNPA address used for originating each multicast packet depends on whether there is a multicast capable IS attached to the subnetworks. When a multicast capable IS is attached, the decision depends on whether there is multicast address mapping information available for that subnetwork corresponding to the group Network address used as the destination address parameter of the multicast packet. When there is a multicast capable IS attached to a subnetwork and there is multicast address mapping information available corresponding to the group Network address, then the SNPA address obtained from the multicast address mapping information is used. Originating multicast packets using the destination SNPA address used for receiving such multicast packets ensures that the multicast packets will not require additional forwarding on the originating subnetwork(s). When there is a multicast capable IS attached to a subnetwork but for which there is no multicast address mapping information available corresponding to the the group Network address, then the SNPA address used is the "All Multicast Capable Intermediate Systems" address.

When there is no multicast capable IS attached to a subnetwork then the ES originating a multicast PDU uses pre-configured information if it is available or the "All Multicast Capable End Systems" SNPA address when no pre-configured information is available.

ES's attached to more than one subnetwork forward each multicast packet that they originate onto every attached subnetwork for which the NSAP address being used as the source address of the multicast packet is actively being reported through the unicast ES-IS Report Configuration function.

6.1.3 Operation of the Active Multicast IS

The Active Multicast IS listens in on all ESGH PDUs originated on the subnetwork for which it is serving as the Active Multicast IS. All subnetworks are handled independently (even if multiple subnetworks have the same ESs attached and the IS is serving as the Active Multicast IS for these subnetworks).

The Active Multicast IS originates MAM PDUs, for all group Network addresses for which it has received ESGH PDUs, on the subnetwork due to the following operational conditions:

- 1) The IS initializes either as the Active Multicast IS after an election with other multicast capable ISs or initializes believing it is the only multicast capable IS;

Note: The determination of such conditions is outside of the scope of this specification;

- 2) The IS receives an ESGH PDU with a group Network address paired to an incorrect SNPA address;
- 3) The expiration of the IS's Multicast Address Mapping Timer for that group Network address; or

Note: This is to prevent the expiration of Mapping Holding Timers in ESs.

- 4) The IS receives a multicast PDU originated on the subnetwork which used an incorrect destination SNPA address.

Note: Of particular concern are those multicast packets using the "All Multicast Capable Intermediate Systems" SNPA address when another SNPA address should have been used. In addition the multicast capable ISs are responsible for listening in on all multicast packets using destination SNPA addresses that are contained within the current multicast address mapping information.

As a result of the event driven conditions (i.e., conditions 2 or 4 above), the Active Multicast IS sends a MAM PDU with direct information (i.e., not needing analysis of the Mask parameters). The Active Multicast IS limits the number of MAM PDUs that are sent out per unit of time. Particular MAM PDUs with direct information will not be sent more than once per second. MAM PDU will be sent in response to continuing event driven conditions such that events occurring greater than 10 seconds after the issuance of such a MAM PDU will result in the issuance of another MAM PDU.

The Active Multicast IS is responsible for forwarding a multicast packet back on the subnetwork it was originated when a multicast packet used the "All Multicast Capable Intermediate System" SNPA address when another SNPA address should have been used. A packet forwarded back onto the subnetwork the multicast packet was originated on will be given a CLNP Lifetime of "1" to prevent the continued relaying of duplicate packets by the multicast ISs.

The further relaying of any multicast packet originated on a subnetwork is the responsibility of the multicast routing protocol used and is outside the scope of this specification.

6.2 Definitions

Active Multicast IS: The one multicast capable IS selected (via means outside of this specification) to originate Multicast Address Mapping information on a particular subnetwork.

Paired SNPA Address: The SNPA address associated with a particular group Network address on a specific subnetwork.

6.3 Routing information supporting multicast transmission

6.3.1 Multicast Announcement Information

An IS should forward a multicast PDU containing a particular destination group Network address onto a subnetwork to which it is attached if and only if one or more of the ESs attached to that subnetwork have declared an interest in receiving multicast PDUs destined for that group Network address. Multicast announcement information enables an IS that supports CLNP multicast to dynamically discover, for each subnetwork to which it is attached, the group Network addresses for which ESs attached to that subnetwork have declared an interest.

On a point-to-point subnetwork the multicast announcement information informs the Network entity, in the case where it is attached to an End System, of the group Network addresses for which that End System expects to receive multicast PDUs.

On a broadcast subnetwork the multicast announcement information informs the multicast capable Intermediate Systems, of the group Network addresses for which ESs attached to that subnetwork expect to receive multicast PDUs.

Note: Intermediate Systems with the optional OSI multicast capabilities do receive information identifying the SNPA address of ESs on the broadcast network that want PDUs with particular group Network addresses as their destination address; however, the critical information is which multicast PDUs are needed, not which ESs need them.

6.3.2 Multicast Address Mapping Information

In order to receive multicast packets destined for a particular group Network address, an ES may need to associate with the group Network address a specific SNPA address. Multicast address mapping information enables an IS to inform ESs that they can receive multicast packets destined for a particular group Network address on a corresponding specific SNPA address. In addition, multicast address mapping information may provide the specific destination SNPA addresses needed by an ES for originating multicast packets.

Multicast address mapping information is not employed on point-to-point subnetworks.

Multicast address mapping information is employed on broadcast subnetworks to enable multicast capable Intermediate Systems to inform the multicast capable End Systems that they can receive, on a specific broadcast subnetwork, multicast packets destined for a particular group Network address on a corresponding specific SNPA address. In addition multicast address mapping information provides the specific destination SNPA address, that corresponds to a particular group Network address, for each multicast packet that it originates on a specific broadcast subnetwork.

6.4 Addresses

All exchanges using this protocol are accomplished over a single subnetwork. While the control PDU's contain Network addresses (i.e., group Network addresses) actual control PDU transfer is accomplished via Subnetwork based group addresses (i.e., group SNPA addresses). The following group SNPA addresses are used: (1)All Multicast Capable End Systems; (2)All Multicast Announcements; (3)All Multicast Capable Intermediate Systems and (4)a group SNPA address corresponding to a group Network address

6.5 Timers

Two additional timers are employed: (1)the Multicast Announcement Timer (MAT) and (2)Multicast Address Mapping Timer (MAMT). Old multicast announcement or multicast address mapping information shall be discarded after the Holding Timer expires to ensure the correct operation of the protocol.

6.5.1 Multicast Announcement Timer

The Multicast Announcement Timer is a local timer (i.e., maintained independently by each End System, one timer per group Network address) which assists in performing the Report Multicast Announcement function. The timer determines how often an End System reports its desire to receive multicast PDUs with that group Network address as its destination address parameter. Considerations in setting this timer are similar to those described for the Configuration timer in the ES-IS specification.

6.5.2 Multicast Address Mapping Timer

The Multicast Address Mapping Timer is a local timer (i.e., maintained independently by an Intermediate System which is actively participating with End Systems to transfer multicast PDUs) which assists in performing the Report Multicast Address Mapping function. The timer determines how often an Intermediate System, actively participating with End Systems for the transfer of multicast PDUs,

reports the Multicast Address Mapping for a particular group Network address. The shorter the Multicast Address Mapping Timer, the more quickly End Systems on the subnetwork will become aware of the correct address mapping which may change due to the Intermediate System becoming available or unavailable. There is a trade off between increased responsiveness and increased use of resources in the subnetwork and in the End Systems.

6.6 Extensions to the current protocol functions

In order to support multicast transmissions the following optional ES-IS protocol functions will be implemented:

6.6.1 Report Configuration by Intermediate Systems

All multicast capable Intermediate Systems on a subnetwork shall use the Multicast Capable option in all ISH PDUs that they originate. This will provide multicast capable End Systems with a way to determine that a multicast capable Intermediate System is operating on a particular subnetwork.

6.6.2 Query Configuration

Note: The Query Configuration function cannot be performed to find the corresponding SNPA address of a group Network 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. While it is possible to design a Query Configuration for use with multicast, this function does not appear to be required given the use of the "All Multicast Capable End Systems" address for supplying a SNPA address when the group SNPA address is not known.

6.7 Multicast Announcement

6.7.1 Report Multicast Announcement Function by End Systems

An End System which needs to receive or continue to receive any multicast PDUs (i.e., PDUs with group Network addresses as their destination address), constructs and transmits ESGH PDUs to inform multicast capable Intermediate Systems of the set of group Network address destinations for which it wishes to receive PDUs. This may be done by constructing ESGH PDUs for each group Network address. Alternatively, ESGH PDUs may be constructed which convey information about more than one group Network 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 Multicast Announcements"
```

If an End System is attached to more than one subnetwork, the information about each group Network address desired for receiving on a particular subnetwork serving the End System shall be transmitted via that subnetwork. It is permissible for an End System to report group Network addresses on multiple subnetworks; however, duplicate multicast PDUs should be anticipated.

The Group Address Pair parameter carries a list of Group Network Addresses, each paired with its associated SNPA address. This information is used by the Active Multicast IS to determine whether a Multicast Address Mapping PDU should be emitted to update the association between Group Network Addresses and SNPA addresses.

The Holding Time (HT) field is set to approximately twice the ES's Multicast Announcement Timer (MAT) parameter. The value shall be large enough so that even if every other ESGH PDU is discarded (due to lack of resources), or otherwise lost in the subnetwork, the multicast announcement information will still be maintained. The value should be set small enough so that Intermediate Systems resources are not needlessly consumed when the ES no longer wishes to receive PDUs destined to a group Network address.

Note: When combining multiple group Network addresses in a single ESGH PDU, it should be realized that there is a single Holding Time parameter associated with all of these addresses.

6.7.1.1 Generating Jitter on Multicast Announcement Timers

The ES shall apply a 25% jitter to its Multicast Announcement Timer (MAT) parameter. When ESGH 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;
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 <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, SendHelloPDU);" where "HelloTime" is 10 seconds, will cause the action "SendHelloPDU" to be performed at random intervals of between 7.5 and 10 seconds. 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 10ms. 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.

Note: Applying jitter to the MAT parameter is required in order to support the optional Damping function. If no jitter is applied on a subnetwork where many ESSs are requesting a particular multicast PDU it is likely that they will have the same value for their MAT and these timers may all become synchronised. Such synchronisation will result in peaks in the distribution of traffic as described above. The resulting overloading of the transmission medium and the systems receiving the PDUs will negate any beneficial use of the Damping function (since systems may be attempting to transmit their own ESGH PDUs at the time they receive ESGH PDUs originated by other ESSs with the same group Network address).

6.7.2 Record Multicast Announcement Function

The Record Multicast Announcement function receives ESGH PDUs, extracts the multicast announcement information and updates the information in its routing information base.

The receiving system is not required to process any option fields in a received ESGH PDU.

Note: When a system chooses to process these optional fields, the precise actions are not specified by this International Standard.

6.7.2.1 Record Multicast Announcement Function by Intermediate Systems

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

The Active Multicast IS upon receipt of an ESGH PDU also extracts the Paired SNPA Address parameter corresponding to each group Network address in the ESGH PDU. If the Active Multicast IS has a mapping for a group Network address carried in the ESGH for which the paired SNPA address does not match, the Report Multicast Address Mapping function is performed.

6.7.2.2 Optional Damping Function

An ES with the optional capabilities to support multicast transfer may decide to process ESGH PDUs multicast by other End Systems. There is potentially some reduction in network traffic by doing this. An ES requesting to receive multicast PDUs is permitted to reset its Multicast Announcement Timer corresponding to one group Network address on one subnetwork upon receiving an ESGH PDU from another ES under the following circumstances:

- a) The {group Network address, paired SNPA address} received on a particular subnetwork matches that of the ES processing the ESGH PDU for that subnetwork.
- b) The Holding Timer parameter value in the ESGH PDU received is equal to or greater than the Holding Timer value for the, group Network address, being used by the ES processing this PDU.

6.7.3 Flush Old Multicast Announcement Function

The Flush Old Multicast Announcement function is executed to remove multicast announcement entries in its routing information base whose Holding Timer has expired. When the Holding Timer for a group Network address expires, this function removes the corresponding entry from the routing information base of the local IS for the corresponding subnetwork.

6.8 Multicast Address Mapping

6.8.1 Report Multicast Address Mapping Function by Intermediate Systems

The Active Multicast Intermediate System constructs a MAM PDU, corresponding to a group Network address for which it received via the Record Multicast Announcement function, and issues these PDUs under the following circumstances:

- a) The IS initializes either as the Active Multicast IS after an election with other multicast capable ISs or initializes after determining it is the only multicast capable IS (the determination of such conditions are outside of the scope of this standard), or

- b) The IS receives an ESGH PDU with a group Network address paired to an SNPA address other than the SNPA address contained in the Active Multicast IS's multicast address mapping information for that group Network address, or

Note: The Active Multicast IS determines which mappings are correct. Pre-configured mappings which are used prior to the initialization of the Active Multicast IS may be determined to be incorrect by the Active Multicast IS.

- c) The expiration of the IS's Multicast Address Mapping Timer for that group Network address.

Note: This is to prevent the expiration of Holding Timers in ESs.

- d) The IS receives a multicast PDU originated on the subnetwork which used an incorrect destination SNPA address.

Note: Of particular concern are those multicast packets using the "All Multicast Capable Intermediate Systems" SNPA address when another SNPA address should have been used. The Originating Subnetwork Forwarding function is performed if this event occurs (see section 6.11).

Note: The multicast capable ISs need to receive multicast packets on all SNPA addresses that are contained in the current multicast address mapping information for the subnetwork. The multicast capable ISs are not required to receive multicast packets on any SNPA addresses other than those contained in the current multicast address mapping information and the "All Multicast Capable Intermediate Systems" SNPA address.

Circumstances b) and d) are the event driven conditions for the Active Multicast IS to construct and issue a MAM PDU. The Active Multicast IS shall limit the number of MAM PDUs issued per unit of time. MAM PDUs with identical information shall not be issued more than once per second. Event conditions occurring 10 seconds after the last issue of an appropriate MAM PDU shall result in the issuance of another such MAM PDU.

The IS serving as the Active Multicast Intermediate System may construct a MAM PDU for each group Network address. Alternatively, MAM PDUs may be constructed which convey information about more than one group Network address at a time, up to the limits imposed by the permitted SNSDU size and the maximum header size of the MAM PDU. The IS performs all multicast address mapping functions independently for each of its subnetworks even if this IS is the Active Multicast IS on multiple subnetworks. Each MAM PDU is transmitted by issuing an SN-

UNITDATA.Request with the following parameters:

```
SN_Userdata (SNSDU) <- MAM PDU
```

```
SN_Destination _Address <- multi-destination address that indicates  
"All Multicast Capable End Systems"
```

The Holding Time (HT) field is set to approximately twice the Intermediate System's Multicast Address Mapping Timer (MAMT) parameter. This variable shall be set to a value large enough so that even if every other MAM PDU, for a particular group Network address, is discarded (due to lack of resources), or otherwise lost in the subnetwork, the multicast address mapping information will still be maintained. The value should be set small enough so that End Systems will quickly cease to use the multicast address mappings supplied by ISs that have failed.

Note: -- The Holding Timer parameter value applies to all group Network addresses called out in the MAM PDU.

The Group Address Pair parameter is used to convey the association between Group Network Addresses and SNPA addresses.

Optionally, the Active Multicast IS may include information in the MAM PDU indicating a larger population of group Network addresses to which the same multicast address mapping information applies. There are two optional fields for this purpose: the Group Network Address Mask option and the Paired SNPA Address Mask option.

There are three permitted cases for including or excluding the masks. In the first case, both masks are absent. In this case the MAM PDU conveys information about one set of enumerated group Network addresses only.

Note: -- Multiple group address pairs may be contained in a single MAM PDU.

In the second case, the MAM PDU contains a Group Network Address Mask but no Paired SNPA Address Mask. In this case, the MAM PDU conveys information about an equivalence class of group Network addresses. The information reveals that multiple group Network addresses are mapped to the same SNPA address.

In the third case, the MAM PDU contains both masks. As in the second case, the MAM PDU conveys information about an equivalence class of group Network addresses. But in this case, the information reveals that the SNPA addresses for the equivalence class of group Network address are embedded in the group Network address. In particular the

Paired SNPA Address Mask indicates the location of the SNPA address in the group Network Address(es).

The Active Multicast IS shall construct a MAM PDU with direct information, not needing analysis of the Mask parameters, in response to the occurrence of an event driven condition. The Active Multicast IS may provide additional information in such a MAM PDU via the use of Mask parameters.

An IS may suggest a value for End Systems on the local subnetwork to use as their Multicast Announcement Timers, for a specific group Network address, by including the Suggested ES Multicast Announcement Timer (ESMAT) parameter in the transmitted MAM PDU. Setting this parameter permits the Active Multicast IS to influence the frequency with which ESs transmit ESGH PDUs.

Note: If the ESMAT parameter is used, the one value permitted in the MAM PDU is suggested for all group Network addresses called out in the MAM PDU.

6.8.2 Record Multicast Address Mapping Function by End Systems

The Record Multicast Address Mapping function receives MAM PDUs, extracts the multicast address mapping information and updates the information in its routing information base. The receiving system is not required to process any option fields in a received MAM PDU with the exception of the Suggested ES Multicast Announcement Timer (ESMAT) parameter.

Note: When a system chooses to process these optional fields, the precise actions are not specified by this International Standard.

On receipt of a MAM PDU an ES with the optional multicast capabilities extracts the multicast address mapping information and stores the {group Network address, paired SNPA address} for a particular subnetwork in its routing information base replacing any other information for the same group Network address and subnetwork.

In addition, an ES shall set its Multicast Announcement Timer, corresponding to the group Network address for which it is performing the Record Multicast Address Mapping function, based on receipt of a MAM PDU, corresponding to that group Network address, containing an ESMAT parameter.

Note: While an ES may process ESGH PDUs multicast by other ESs to support the optional Damping function, an ES is not permitted to change its own mapping due to the mapping found in other ES's ESGH PDUs.

6.8.3 Flush Old Multicast Address Mapping Function by End Systems

The Flush Old Multicast Address Mapping function is executed to remove multicast address mapping entries in its routing information base whose corresponding Holding Timer has expired. When such a Holding Timer for a multicast address mapping expires, this function removes the corresponding entry from its routing information base for the corresponding SNPA.

6.9 Paired SNPA Address Selection Function by End Systems

An End System shall pair each group Network address with an associated SNPA address to support receiving (e.g., performing the Report Multicast Announcement function) and originating multicast PDUs.

6.9.1 Paired SNPA Address Selection for Receiving Multicast PDUs

An End System always has a paired SNPA address for every active group Network address on a particular subnetwork. This mapping is obtained by:

- a) recording a multicast address mapping which is maintaining an active holding timer, or if there has been no dynamic information received, by
- b) having pre-configured multicast address mapping information, or if neither dynamic nor pre-configured information is available, by
- c) mapping the "All Multicast Capable End Systems" multi-destination address to the group Network address.

6.9.2 Paired SNPA Address Selection for Originating Multicast PDUs

An End System, originating a multicast PDU, pairs a SNPA address to the group Network address. This mapping is obtained in the following manner:

- a) If there is a multicast capable IS reachable on the subnetwork then the SNPA address used by an End System originating a multicast PDU is either the paired SNPA address obtained from the multicast address mapping information associated with the group Network address in the multicast PDU's Destination address parameter or if there is no valid entry for the group Network address by using the "All Multicast Capable Intermediate Systems" multi-destination address, or if there is no multicast capable Intermediate System on the subnetwork, by

Note: Multicast address mapping information is valid if the Holding Timer associated with it has not expired.

Note: An ES can determine if a multicast capable IS is reachable on the subnetwork by having for that subnetwork either (1) multicast address mapping information or (2) routing information received via an ISH PDU containing a Multicast Capable optional parameter. In either case the information must be valid (i.e., the Holding Timer for the information must not have expired).

- b) having pre-configured multicast address mapping information, or if neither a multicast capable Intermediate System is present on the subnetwork nor pre-configured information is available, by
- c) mapping the "All Multicast Capable End Systems" multi-destination address to the group Network address.

6.10 Extensions to the ISO CLNP Route Function by End Systems

An End System attached to more than one subnetwork shall determine when originating a multicast PDU whether to forward this multicast PDU to more than one subnetwork or not. End Systems shall originate each multicast PDU on all subnetworks for which the ISO ES-IS Configuration function is actively reporting the NSAP address contained in the Source Address parameter of the multicast PDU. As a result of this function multiple invocations of the ISO CLNP Forwarding function may result when such an ES originates a multicast PDU.

6.11 Originating Subnetwork Forwarding Function by Intermediate Systems

The Active Multicast IS upon receiving a multicast PDU originated on a subnetwork which used the "All Multicast Capable Intermediate Systems" SNPA address when another SNPA address should have been used, performs the Originating Subnetwork Forwarding function. The multicast address mapping information defines the correct SNPA address pairings for a given subnetwork. The Originating Subnetwork Forwarding function forwards the multicast PDU back on subnetwork it was originated on. In the case that the ES was attached to more than one subnetwork and originated the multicast PDU on more than one subnetwork, the Active Multicast IS for each subnetwork performs the Originating Subnetwork Forwarding function for the subnetwork that they are responsible for.

The Active Multicast IS obtains the contents for the multicast PDU for the Originating Subnetwork Forwarding function by using the contents of the multicast PDU received with the incorrect destination

SNPA address and replacing the original PDU Lifetime field with the value one (0000 0001). The Active Multicast IS performs the ISO 8473 PDU Composition function and forwards the PDU to the subnetwork that the PDU was originated on using the ISO 8473 Forwarding function with the correct destination SNPA address.

Note: The PDU Lifetime field is set to "one" to ensure that ISs attached to the originating subnetwork do not forward this PDU on. Such ISs should have received the PDU when it was originated since this function is only performed in the event of receiving a multicast PDU incorrectly addressed to the "All Multicast Capable Intermediate Systems" SNPA address.

6.12 Structure and Encoding of PDUs

The ES-IS multicast control functions are supported via the exchange of ESGH and MAM PDUs. The one exception to this is that a new optional parameter, the Multicast Capable parameter, is provided for use within the ISH PDU.

6.12.1 PDU Type Codes

The Multicast Announcement is accomplished via the transfer of End System Group Hello (ESGH) PDUs. The PDU type code for an ESGH PDU is "0 0 1 0 1". The Multicast Address Mapping (MAM) is accomplished via the transfer of Multicast Address Mapping PDUs. The PDU type code for a MAM PDU is "0 0 1 1 1".

6.12.2 Hold Time field

The Holding Time field specifies the maximum time for the receiving Network entity to retain the multicast announcement or multicast address mapping information contained in the PDU.

6.12.3 Structure of Addressing Parameters

The ESGH and MAM PDUs carry one or more group Network addresses (GNAs) each with their associated Paired SNPA Address (PSA).

6.12.4 Group Address Pair Parameter for ESGH and MAM PDUs

The Group Address Pair parameter is a list of one or more group Network addresses each with their associated Paired SNPA address. The group Network address identifies specific multicast PDUs and the Paired SNPA address is the SNPA address on which the ES expects to receive such multicast PDUs on that subnetwork. It is encoded in the ESGH and MAM PDUs as shown in Figure 1.

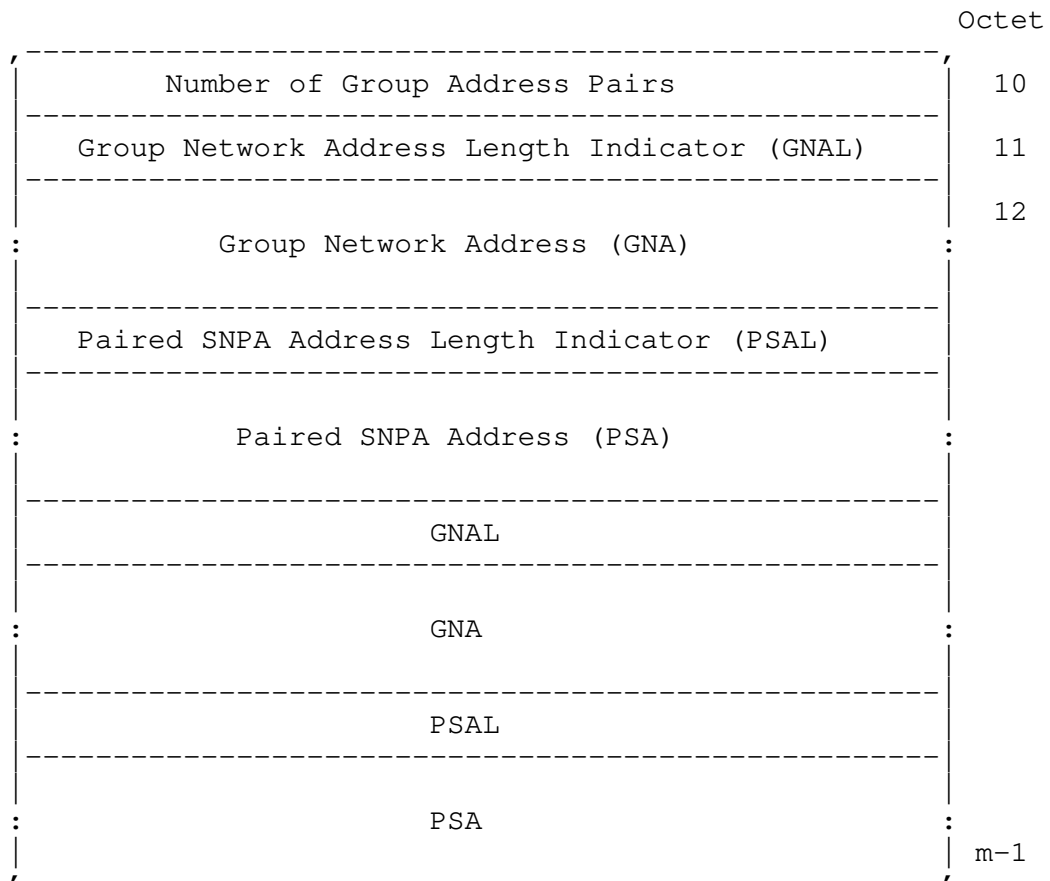


Figure 1 - ESGH and MAM PDUs - - Group Address Pair Parameter

6.12.5 Extensions to the current Option Parameters

The Security and Priority optional parameters may be carried in a ESGH PDU. There is no Security or Priority option for the MAM PDU.

6.12.6 Suggested ES Multicast Announcement Timer

The ESMAT parameter may appear only in the MAM PDU

The ESMAT parameter conveys the value that an IS requests the receiving ESs to use as their local Multicast Announcement Timer.

Parameter Code: 1100 0111

Parameter Length: two octets

Parameter Value: ESMAT in units of seconds.

6.12.7 Multicast Capable

The Multicast Capable option may appear only in the ISH PDU

The Multicast Capable options consists only of a one octet code and a one octet parameter length field, there is no parameter field.

Parameter Code: 1100 1000

Parameter Length: zero octets

Parameter Value: none (parameter does not exist).

6.12.8 Group Network Address Mask

The Group Network Address Mask option may only appear in the MAM PDU.

The Group Network Address Mask parameter indicates that the multicast address mapping information applies to a larger population of group Network Addresses than the group Network address(es) contained in the MAM PDU indicates. When this option is provided in a MAM PDU, the masking relationship contained must be valid for all group Network addresses contained in this PDU. An End System may ignore this parameter.

The Group Network Address Mask establishes an equivalence class of group Network addresses to which the same multicast address mapping information applies. To determine whether or not a trial group Network address falls within the equivalence class, the ES aligns the trial group Network address with the Group Network Address Mask padding the latter with trailing zero octets if necessary. If in all bit positions where the Group Network Address Mask is "1" the trial group Network address matches the Group Network Address field of the Group Address Pair parameter of the MAM PDU, then the trial group Network address belongs to the equivalence class described by the MAM PDU.

The Group Network Address Mask parameter has additional semantics when considered with the Paired SNPA Address Mask parameter.

Parameter Code: 1110 0011

Parameter Length: variable, up to 20 octets

Parameter Value: a comparison mask of octets to be aligned with the Group Network Address field of the Group Address Pair parameter of the MAM PDU.

6.12.9 Paired SNPA Address Mask

The Paired SNPA Address Mask option may only appear in the MAM PDU.

When the Paired SNPA Address Mask is present, the equivalence class defined by the Group Network Address Mask also has common structure below the Group Network Address Mask; i.e., in the portion of the group Network address where the Group Network Address Mask is logically "0". The Paired SNPA Address Mask supplies additional information about the structure, by indicating certain bit positions within the space "below" the Group Network Address Mask. Specifically, the Paired SNPA Address Mask indicates the location of the Paired SNPA address in the Group Network Address.

This parameter may appear in a MAM PDU only if the Group Network Address Mask is also present. When this option is provided in a MAM PDU, the masking relationship contained must be valid for all group Network addresses contained in this PDU. An ES receiving such a MAM PDU may safely ignore both masks. However (since presence of both masks dictates different functional behavior than the presence of the Group Network Address Mask alone) an ES shall not ignore one of the masks while heeding the other.

Parameter Code: 1110 0100

Parameter Length: variable

Parameter Value: a comparison mask of octets to be aligned with the Group Network Address field(s) of the Group Address Pair parameter of the MAM PDU.

6.12.9.1 Mask Parameters Example

This section provides examples of using the Group Network Address Mask and the Paired SNPA Address Mask. The examples given are for an Internet usage of CLNP Multicasting across subnetworks using IEEE 802 addressing. For these examples the group Network address format is:

| | |
|---------|---|
| | +-----+-----+-----+-----+-----+-----+-----+-----+ |
| | IDP Upper DSP Embedded SNPA address SEL |
| | +-----+-----+-----+-----+-----+-----+-----+-----+ |
| octets: | 3 10 6 1 |
| | +-----+-----+-----+-----+-----+-----+-----+-----+ |

Thus the group Network address used is 20 octets. For these examples, the only field considered is the Embedded SNPA address field and its placement within the group Network address.

In the first example it is the policy in "this part of the Internet" to map the Embedded SNPA address into the IEEE 802 address space reserved by IEEE 802 for group addressing using LOCAL assignment, this corresponds to all 48 bit values with the two low order bits of the first octet set to "11".

The Active Multicast Intermediate System on this subnetwork may construct a MAM PDU to map, for this example, a group Network address of {13 octets, 03-00-DA-DA-DA-DA, 1 octet} and a paired SNPA address of 03-00-DA-DA-DA-DA. In addition the Active Multicast Intermediate System can include in the MAM PDU a Group Network Address Mask of FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-03-00-00-00-00-00-00.

With this parameter, all group Network addresses which share the identical first 13 octet and with "11" in the two low order bits of the 14th octet are put in an equivalence class and share the same mapping information. If this were the only option present then all of these group Network addresses would all have a paired SNPA address of 03-00-DA-DA-DA-DA.

In order to map the group Network addresses to the range of IEEE addresses of this example, the MAM PDU must also contain a Paired SNPA Address Mask. The Paired SNPA Address Mask identifies where the SNPA Address is contained within the group Network addresses (defined by the equivalence class formed by the Group Network Address Mask within the same PDU). For this example the Paired SNPA Address Mask is 00-00-00-00-00-00-00-00-00-00-00-00-00-00-FF-FF-FF-FF-FF-FF-00.

As a second example, all group Network addresses with a specific OUI (organizationally unique identifier) using the twenty octet group Network address format provided above are mapped to their embedded SNPA address. An OUI is assigned by IEEE 802 and is three octets in length. The OUI is contained in the first three address octets of a GLOBALLY assigned IEEE 802 address. For this example the MAM PDU must contain the following:

1. A group Network address contained within the MAM PDU with the OUI of interest.
2. A group Network address Mask of FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-FF-00-00-00-00.
3. A Paired SNPA Address of 00-00-00-00-00-00-00-00-00-00-00-00-00-00-FF-FF-FF-FF-FF-FF-00.

6.12.10 End System Group Hello (ESGH) PDU

The ESGH PDU has the format shown in figure 2:

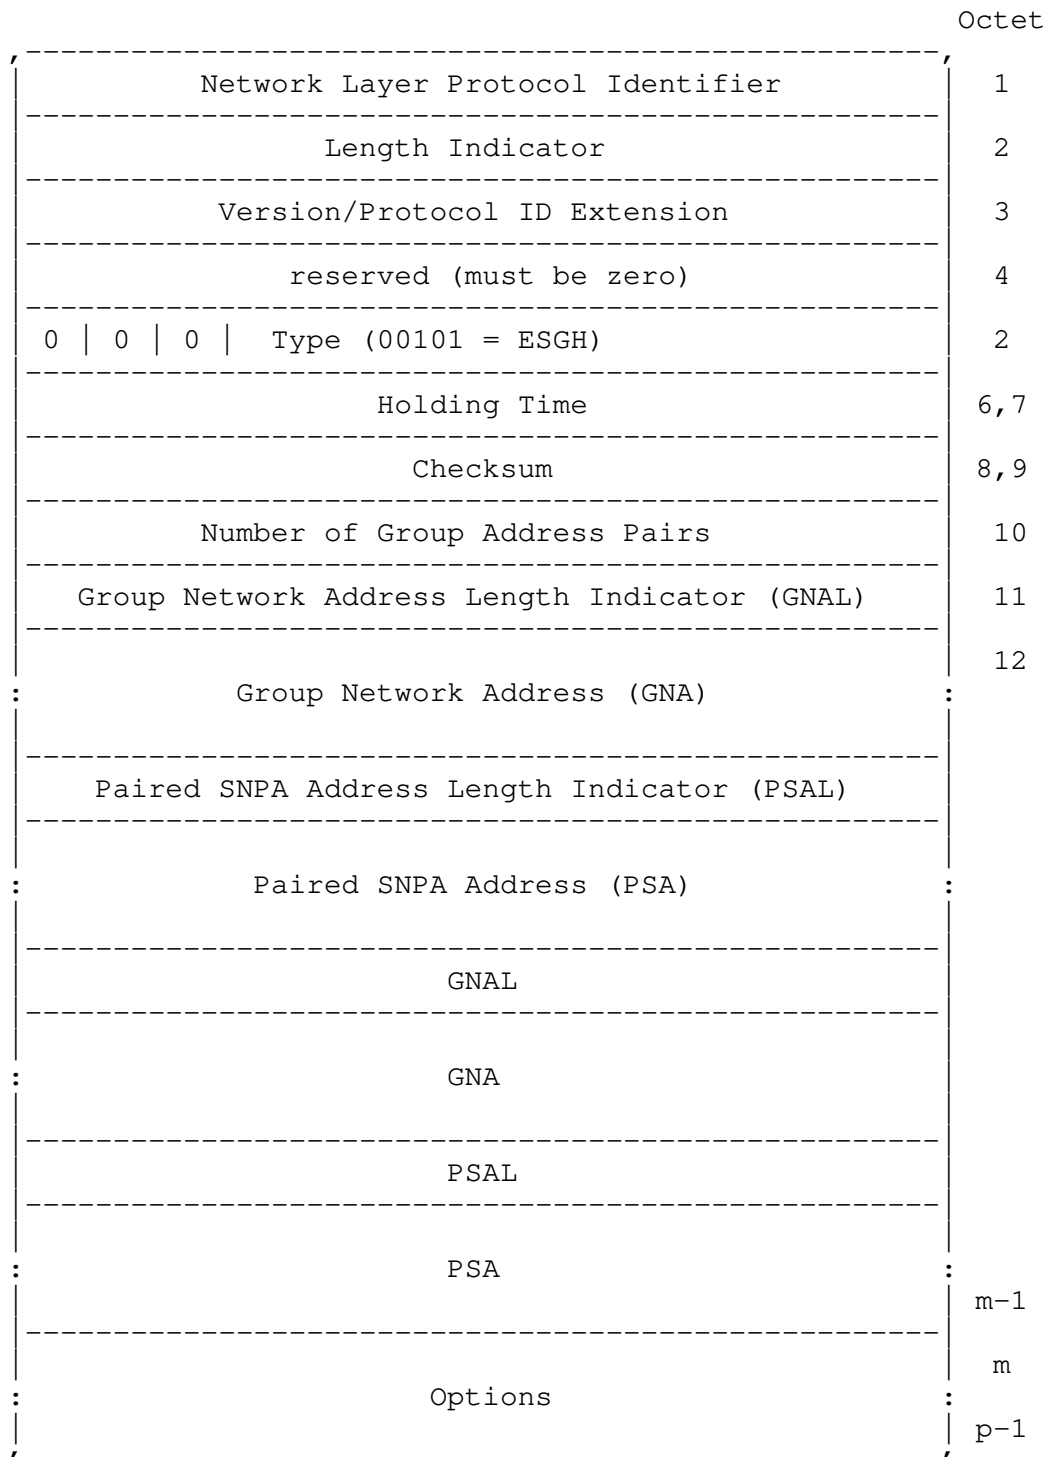
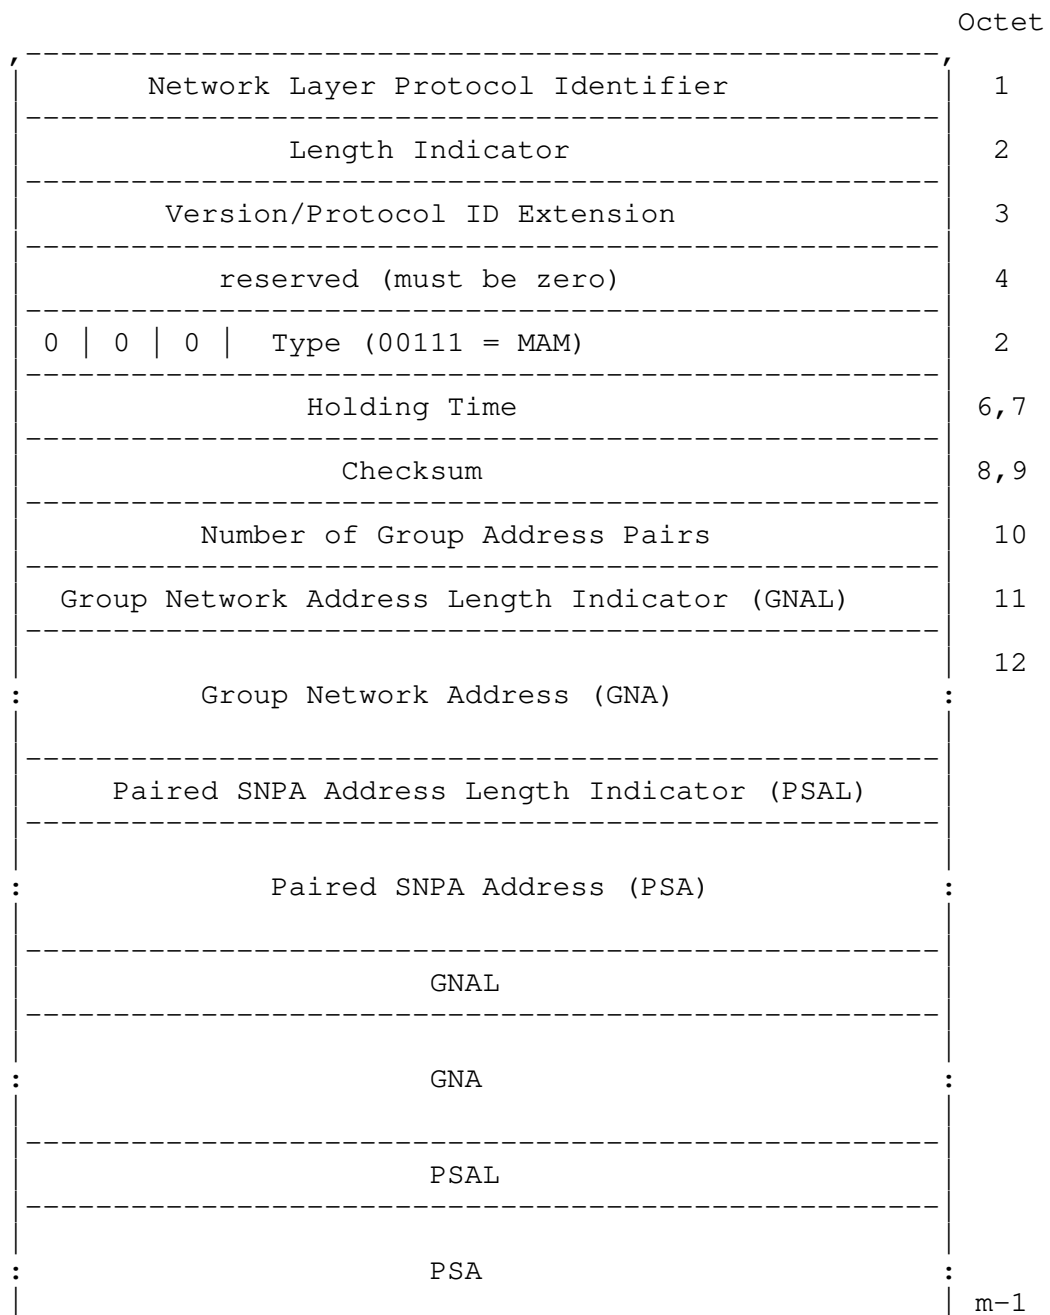


Figure 2 - ESGH PDU Format

6.12.11 Multicast Address Mapping (MAM) PDU

The MAM PDU has the format shown in figure 3:






Figure 3 - MAM PDU Format

6.13 Conformance

All of the extensions provided to the functions to support multicast capability are optional. For an End System or Intermediate System which is not multicast capable these extensions are not applicable. A Network entity may choose to be multicast capable, a multicast capable Network entity is required to support both multicast announcement information and multicast address mapping information.

An implementation claiming conformance as a multicast capable End System shall meet all of the requirements for an End System which is not multicast capable and shall support multicast announcement information and shall implement the functions marked as Mandatory (M) in column 4 of table 3. A multicast capable End System implementation shall also support multicast address mapping information and shall implement the functions marked as Mandatory (M) in column 5 of table 3.

An implementation claiming conformance as a multicast capable Intermediate System shall meet all of the requirements for an Intermediate System which is not multicast capable and shall support multicast announcement information and shall implement the functions marked as Mandatory (M) in column 6 of table 3. A multicast capable Intermediate System implementation shall also support multicast address mapping information and shall implement the functions marked as Mandatory (M) in column 7 of table 3.

Table 3 - Static Conformance Requirements for Multicast Capable Network Entities

| Label | Function | Clause Reference | ES | | IS | |
|---------|-----------------------------------|---------------------|----|----|----|----|
| | | | AI | MI | AI | MI |
| RpMAN | Report Multicast Announcement | 6.7.1 | M | - | - | - |
| RcMAN | Record Multicast Announcement | 6.7.2.1 | - | - | M | - |
| RcDamp | Record Damping | 6.7.2.2 | O | - | - | - |
| FlMAN | Flush Old Multicast Announcement | 6.7.3 | O | - | M | - |
| RpMAdMa | Report Multicast Address Mapping | 6.8.1 | - | - | - | M |
| MATGn | ESMAT Generation | 6.8.1 | - | - | - | M |
| RcMAdMa | Record Multicast Address Mapping | 6.8.2 | - | M | - | - |
| MATPr | ESMAT Processing | 6.8.2 | - | M | - | - |
| FlMAdMa | Flush Old Multicast Address Map | 6.8.3 | - | M | - | - |
| PSAdSel | Paired SNPA Address Selection | 6.9.1 | - | M | - | - |
| ExtForw | Extensions to CLNP Route Function | 6.10 | - | M | - | - |
| OSuForw | Originating Subnetwork Forwarding | 6.11 | - | - | - | M |

Key:

AI = Multicast Announcement information supported

MI = Multicast Address Mapping information supported

M = Mandatory; O = Optional; - = not applicable

7. Security Considerations

Security issues are not discussed in this memo.

Appendix A. Differences with RFC 1112

This appendix is intended to identify differences between the mechanisms defined for CLNP Multicast in this specification and those for IP multicast defined in RFC 1112. The work on CLNP Multicast followed the work on IP multicast and was explicitly aimed at bringing the capabilities described in RFC 1112 into a CLNP context. This appendix is intended to provide some background information on the difference; however, it is not intended to justify the mechanisms selected for CLNP multicast use.

Static/Dynamic Address Binding of Multicast Datagrams

IP multicast utilizes a static binding of Class D IP addresses to a specific range of IEEE 802 48 bit group addresses. The IEEE 802 address range that is used is within the address range that IEEE 802 allocates for "Global" administration and this block of addresses is under the control of the Internet Assigned Numbers Authority (IANA) which in turn has allocated this block of addresses for use by IP multicast. This scheme is very simple and efficient. Given the use of a 32 bit IP address, the lower 23 bits of the Class D address are mapped into the lower 23 bits of a 48 bit IEEE 802 address where the upper 25 bits are fixed. Static binding of this form is global in scope (all members of a group use the same IEEE 802 address on all subnets (at least all that use IEEE 802 addressing)).

CLNP multicast uses a dynamic binding of a group Network address (up to 20 bytes) to any subnetwork address. In cases where no multicast capable Intermediate Systems are attached to a subnetwork then a binding using preconfigured information or the "All Multicast Capable End Systems" subnetwork addresses is used. The large GNA provides the room to contain a full 48 bit IEEE 802 address if desired. Mask capabilities are optionally provided which allow a multicast capable Intermediate System to specify a "static" binding for a particular subnetwork. One of the major purposes of providing a dynamic binding is to customize a host's subnetwork address usage to the capabilities of the attached systems. There is considerable differences in the numbers of group subnetwork addresses that a system can recognize using hardware hooks built into the integrated circuits used. For example the number of addresses that can be recognized by hardware may differ by an attached system depending upon the interface it uses (e.g., Ethernet interface and FDDI within the same system may have quite different capabilities). Dynamic binding of this form is local in scope (members of a group may use different subnetwork addresses (e.g., IEEE 802 addresses) on different subnets).

Originating of Multicast Datagrams

IP multicast originates multicast datagrams directly, where the host originating a datagram sends it with the group Subnetwork address as its destination. Hosts attached to the network where the datagram is originated receive the datagram directly.

CLNP multicast originates multicast datagrams directly using the group's subnetwork address as its destination when multicast address mapping information is available. This case occurs when a multicast capable Intermediate System is attached to the subnetwork and a host on the subnetwork is announcing an interest in multicast packets identified by a particular group Network address. The Active Multicast IS may use MAM PDU mask parameters to provide multicast address mapping information for a large number of group Network addresses. When there is no multicast address mapping information for the particular group Network address on a subnetwork with a multicast capable IS attached to it, hosts originate packets using such addresses sends to the "All Multicast Capable Intermediate Systems" SNPA address. This case occurs when there are no receivers of such multicast packets on the originating subnetwork. When a multicast capable Intermediate System is not attached to a subnetwork, the End System may utilize either preconfigured information (which might be a direct mapping from a portion of the group Network address) or use the "All Multicast Capable End Systems" address.

Address Binding of Control Packets

IP multicast sends the control packets related to the IGMP protocol on the same subnetwork address that is used by the multicast data traffic.

CLNP multicast sends the control packets related to the ES-IS protocol extensions on specific group subnetwork addresses (i.e., "All Multicast Capable End Systems" and "All Multicast Announcements" addresses).

Router Requirements for relaying Multicast Datagrams

IP multicast requires that a multicast router run in "promiscuous" mode where it must receive all multicast datagrams originated on a subnetwork regardless of the destination. This is a result of the choices selected in the "Originating of Multicast Datagrams" and "Address Binding of Control Packets" discussed above.

CLNP multicast allows a multicast router to limit multicast packet reception to only those datagrams sent to the SNPA addresses where there is current multicast address mapping information or to the "All

Multicast Capable Intermediate Systems" address. The intention is to allow the multicast routers to be in control of the SNPA addresses for multicast packets that they need to receive. This is a result of the choices selected in the "Originating of Multicast Datagrams" and "Address Binding of Control Packets" discussed above.

Aggregation of Control Information

In IP multicast, a host is required to withhold an announcement report upon hearing another host reporting a similar interest in a particular Class D address on a particular subnetwork. This is an option for CLNP multicast (upon hearing interest in a particular group Network address on a particular subnetwork). Such reports are not combined in IP multicast while CLNP multicast supports providing multiple announcements (and address mappings) within a single packet. A mask feature for address mappings supports identifying mappings for a range of group Network addresses within a single control packet.

Datagram Scope Control

IP multicast supports the use of the IP Hop Count as a means to support scope control. While not documented in RFC 1112, a technique is also being used to use bits within the Class D address to identify whether a datagram has single subnetwork, "campus" or global scope.

CLNP has considerable scope control functionality. While the PDU Lifetime field can be employed in a similar way to the IP Hop Count, two additional options are available. The Radius scope control provides a mechanism for "administratively" setting distance values and de-couples the multicast scope control from the PDU lifetime function. More importantly, the Prefix based scope control appears to provide considerable and flexible functionality that can adjust to situations where a known, hierarchical unicast addressing structure exists.

Marking of Multicast Datagrams

IP multicast marks a multicast PDU via the use of an IP Class D address as its destination address parameter. CLNP multicast marks both the PDU (a different PDU type) and the destination address (i.e., group Network address) parameter.

Unicast Addressing Differences

An IP address identifies a specific host interface while a CLNP individual Network address (i.e., NSAP address) identifies a particular Network entity. This difference has led to a difference with RFC 1112. IP multicast requires a host which is attached to

more than one subnetwork to originate a multicast packet on only one subnetwork. CLNP multicast requires a host which is attached to more than one subnetwork to originate a multicast packet on every subnetwork that the ISO ES-IS Configuration function is reporting the NSAP address contained in the source address parameter of the multicast PDU.

Error Reports

Error reports sent in response to receiving a multicast PDU are not permitted in IP multicast while they are permitted in CLNP multicast.

Source Routing

Source routing of multicast PDUs are permitted in IP multicast (but at the present time this is discouraged) while they are not permitted in CLNP multicast.

Appendix B. Issues Under Study

This appendix is intended to record the current issues (as discussed at the March 1994 TUBA meeting).

1. Local versus Global address bindings

The extensions to the ES-IS protocol provide a multicast address mapping function which supports dynamically binding a group Network address to a subnetwork address. Concern has been expressed that this is an unnecessary feature which complicates the job of network administrators without suitable benefit. A static, global binding of group Network addresses to IEEE 802 subnetwork addresses, as is used by IP multicast has been suggested.

The two main reasons that the group Network address to subnetwork (IEEE 802) address was made locally configurable were to support multicast on subnets with hosts having a mixture of capabilities (as to how many multicast subnetwork addresses a host could register to receive at a time) and to support multicast on subnets that do not use 48 bit IEEE 802 addresses. Thus it was felt that this should be done per subnetwork versus globally. Even multi-homed hosts with subnets that use 802 addresses may have varying capabilities (looking at typical Ethernet, FDDI and 802.5 implementations).

One possible solution is to recommend a direct mapping in any Internet use of CLNP multicast on subnets which use IEEE 802 addressing. This could be a default for all Internet hosts. A policy would be needed to identify the Internet's group Network address format. Given such a mapping the only operational overhead

that would occur is that in the presence of a mapping server (the Active Multicast IS), which was supporting this mapping, a MAM PDU would periodically be sent with a Group Network Address Mask which would identify the direct mapping.

2. "Real Time" Scope Control Features

The scope control features are provided via optional parameters. Use of multicast transfer of audio and video streams may require scope control mechanisms which operate very quickly.

One possible solution is to embed scope control mechanisms into the group Network address itself. For example, a group Network address using the "Local" AFI is automatically limited to not cross inter-domain borders. Further, more flexible, address formats may be developed.

References

[Deering91] Deering, S., "Multicast Routing in a Datagram Internetwork", PhD thesis, Electrical Engineering Dept., Stanford University, December 1991.

[RFC1112] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, Stanford University, August 1989.

[RFC1237] Colella, R., Gardner, E., and R. Callon, "Guidelines for OSI NSAP Allocation in the Internet", RFC 1237, NIST, Mitre, DEC, July 1991.

[CLNP] Protocol for providing the connectionless-mode network service, International Standard 8473-1, Second Edition, ISO/IEC JTC 1, Switzerland 1994. (Available via FTP from merit.edu:pub/iso/iso8473part1.ps).

[ES-IS] End system to Intermediate system routing exchange protocol for use in conjunction with the Protocol for providing the connectionless-mode network service, International Standard 9542, ISO/IEC JTC 1, Switzerland 1987. (Available via FTP from merit.edu:pub/iso/iso9542.ps).

[MULT-AMDS]: Amendments to ISO standards to support CLNP multicast extensions:

ISO 8348 AM5 Amendment to the Network Service to support Group Network Addressing. International Standard ISO 8348 Amendment 5, ISO/IEC JTC 1, Switzerland 1994.

ISO 8473-1 DAM1 - Draft Amendment to the Second Edition of the Protocol for providing the connectionless-mode network service [CLNP], Multicast Extension, 1993.

ISO 9542 DAM2 - Draft Amendment to the ES-IS [ES-IS] protocol, Addition of connectionless- mode multicast capability, 1993.

Author's Address

Dave Marlow
Code B35
NSWC-DD
Dahlgren, VA. 22448

Phone: (703) 663-1675
EMail: dmarlow@relay.nswc.navy.mil

