ATNP WG 2/**WP 562R1**
ATNP WGW/ **WP 3-14R1**
December 9, 1999

ATNP Working Group of the Whole
3 rd Meeting
Tokyo, Japan
6 December 1999

Transmittal Paper on Multicast Architecture to ATNP/3

Prepared by: James Moulton

Presented by James Moulton

Summary

The attached paper is the draft proposed Working Paper for presentation at ATNP/3 on the subject of the work on Multicast.

ATNP/3-WP__

AERONAUTICAL TELECOMMUNICATION NETWORK PANEL
THIRD MEETING

Montreal 7th - 18th February 2000

Agenda Item 3: Development of SARPs and guidance material relating to ATN systems management and security as well as to the enhancement of existing ATN functionalities and applications

ATN Internetwork Working Group (WG2) - Report on Provisions for Multicast since ATNP/2

(Presented by ATNP WG2 Rapporteur)

WORKING PAPER

SUMMARY

At ATNP/2, WG 2 was assigned the task of investigating the inclusion of multicast/broadcast capabilities within the ATN. WG 2 has progressed the architectural work on multicast but was not able to complete additions to the Doc. 9705 incorporating this functionality. The following paper and attachment provides details about the status of the work on multicast/broadcast capabilities.

REFERENCES

1. ATNP/2 Report on Agenda Item 5
2. Proposed work programme beyond ATNP/3

1 Introduction

1.1 The ATNP/2 approved work plan for the ATN Internet Working Group included as item (h) the “investigation of the provision of multicast/broadcast functions in the ATN;”. During the ten meetings of WG 2, work has progressed on the definition of a basic approach to adding multicast/broadcast functions to the ATN.

1.2 This paper summarizes the progress made by WG 2 in defining multicast/broadcast functions to the ATN and provides information about the future work required to add this functionality.

2 Discussion

2.1 Work Progress

2.1.1 During its deliberations, seven Working Papers were introduced to WG 2 on the subject of multicast architecture and services. These Working Papers provided significant information on approaches for adding multicast/broadcast services in the ATN Internet.

2.1.2 From the analysis of the amount of work required to add multicast/broadcast functions to the ATN Internet, it was concluded that inclusion of the new functionality during the period between ATNP/2 and ATNP/3 was not possible.

2.1.3 The current state of the work, as progressed by WG 2, is attached to this paper as Attachment A. This work is recommended for future work on the addition of multicast/broadcast functions after ATNP/3.

2.2 Overview of Current Work

2.2.1 WG 2 has performed preliminary work in the definition of multicast/broadcast functions for the ATN Internet. The work focused on two important areas: the architecture for multicast/broadcast services, and the applicability of existing standards.

2.2.2 WG 2 analyzed the work performed in ISO and ITU on the definition of a multicast reference model. The work formed the basis of the definition of multicast/broadcast capabilities.

2.2.3 WG 2 analyzed the requirements for adding multicast/broadcast services, while maintaining backwards compatibility. It was recognized early in the work effort that maintaining backwards compatibility and adding new functionality was very difficult, especially in the case of multicast/broadcast where the new functionality is required along the entire communication path. These architectural and deployment issues are central to the successful definition of the new functions.

2.2.4 WG 2 analyzed the ISO work in defining multicast extensions to the ATN Internet protocols. The analysis identified areas that need further refinement and enhancements in order for the standards to be usable in the ATN environment.

3 Recommendation

3.1 The Panel is invited to note the progress made by WG 2 in defining multicast/broadcast features provided as Attachment A and to place the continuation of the work in the future work plan.

ATTACHMENT A - MULTICAST ARCHITECTURE

1 Introduction

1.1 During the course of its work subsequent to ATNP/2, WG 2 has investigated the addition of multicast/broadcast functions to the ATN. A series of Working Papers were presented at the Working Group meetings that presented details on the subject and provided possible solutions.

1.2 This paper presents details on the outcome of the work performed on defining multicast/broadcast functions within the ATN Internet.

2 Background

2.1 Within the ATN Internet, there may be some services which lend themselves to either a multicast or broadcast operation. In fact, the use of a broadcast media such as radio make the implementation of either broadcast or multicast easier and more effective.

2.2 The ISO and ITU-T committees responsible for the OSI Reference Model and the lower layer services have been working on multicast and broadcast architecture, services, and protocols for the last several years.

2.3 Proposals for enhancements to the ATN Internet SARPs were received. Separate proposals were received that were based on the OSI standardized multicast extensions as well as ATN specific extensions.

3 Paper Organization

3.1 Section 4 presents a background on multicast/broadcast architectures.

3.2 Section 5 presents the status of the OSI multicast extensions.

3.3 Section 6 presents an analysis of the different options.

3.4 Section 7 presents future plans.

4 Multicast Architecture

4.1 Overview

4.1.1 Multicast architecture is a complex subject with many different aspects. This section briefly overviews how multicast may be viewed in an ATN-type environment.

4.1.2 Multicast consists of a one-to-many architecture where an application can send data to a (potentially) large number of receivers that specifically request to receive that data.

4.1.3 To transmit data to a group of multicast subscribers, the transmitting end-system sends its data to its nearest multicast-enabled router. This router replicates the packet to other multicast routers within the group. The distribution of multicast data resembles a tree with the originator of the transmission sitting at the top of the tree. Since the transmitting end-system only needs to send one stream of data to the network, there is a benefit of reduced performance requirements of that system as a reduced bandwidth requirements for large amounts of data distribution.

4.1.4 An end-system subscribes to a multicast group and determines the group's group NSAP address. The routing protocols support group address routing and creates the routing trees.

4.2 Detailed Multicast Concepts

4.2.1 Multicasting Groups

4.2.1.1 Group Membership, Transmission, and Addressing

4.2.1.1.1 The concept of groups is central to understanding the operation and variations of multicast.

4.2.1.1.2 A group is a set of recipient entities possibly designated by a group address. There are several different possibilities regarding the relationship between sending entities and groups, and the extent of knowledge that sending and receiving entities need to have about groups in a multicast transmission.

4.2.1.2 Closed Group Membership

4.2.1.2.1 A closed group is a set of recipients that are each only capable of receiving multicast messages originating from other members of the group. Non-members are not permitted to transmit multicast messages to closed groups. Closed groups are useful for multicast applications such as teleconferenceing, in which only conference participants would want to transmit messages to other conference participants.

4.2.1.3 Open Group Membership

4.2.1.3.1 An open group is a set of recipients which are capable of receiving multicast messages from all other peer entities, regardless of whether or not the originator is a member of the group. Because the set of originating entities is unrestricted, there is no way to provide restricted multicast transmission to an open group. That is, there is no way to ensure that only one multicast transmission will be sent to the group at any given moment.

4.2.1.4 Static Group Membership

4.2.1.4.1 Static groups are defined by system management in a process that is outside of the operation of an instance of communication between the entities of the group. Membership in such groups cannot use multicast transmissions to alter membership in the group. Once a static group is registered, transmission of multicast messages to that group are enabled.

4.2.1.5 Dynamic Group Membership

4.2.1.5.1 Dynamic groups have the potential to be constantly changing. Once a dynamic group is registered, transmission of multicast message to that group are enabled. Members can, however. Leave and enter the group at will by communicating their desire to do so. Multicast messages transmitted will be delivered to each of the current members of the group. Dynamic groups require that entities be aware that groups exist and that entities know what groups they belong to so that they can enter and leave groups.

4.2.1.6 Indeterminate Group Membership

4.2.1.6.1 An indeterminate group is a nebulously defined set of entities in that not all indeterminate group members are necessarily aware of the identities of all other group members. Communication destined for indeterminate groups can be transmitted without knowledge of or regard for how many recipients may actually be participating in the communication. Hence, receipt of such data transmissions to all members of such an indeterminate group cannot be ensured. Confirmed service cannot be provided to indeterminate groups because the absence of a given acknowledgement would not necessarily even be noticed. Negatively acknowledged service might be possible, provided that the originator can be made aware of the identity of the destination to which the lost or corrupted data should be retransmitted. Broadcast television is one example of an application in which multicast transmissions are made to an indeterminate group.

4.2.1.7 Determinate Group Membership

4.2.1.7.1 A determinate group is a well-defined set of entities, all the identities of which are well-known and maintained as state information by all group members. Both reliable and unreliable data transmissions can be transmitted among determinate groups, because knowledge of the identities of all group members enables an originator to detect the absence of a given expected acknowledgment. If retransmission is necessary, the identity of the recipient of the retransmission is known to the originator.

4.2.1.8 Fixed, single Transmission Source

4.2.1.8.1 In the case of a group having a fixed, single transmission source, there is only one fixed transmitter that sends multicast messages to a group. Return messages from each of the group members back to the fixed source may or may not be allowed. If they are allowed, such return messages may be either connectionless or connection-oriented. If they are connectionless, then they resemble unicast response to the originator of the multicast. If they are part of a connection-oriented transmission, then the connection between the originator and multiple recipients is essentially a 1-to-n connection. In the connectionless case, the recipients need not necessarily even know that they have received a multicast message or that they are members of a group, let alone who the other members are of the group.

4.2.1.9 Closed Group, Single Transmission Source

4.2.1.9.1 A closed group with restricted transmission is similar to a fixe transmitter regime with the additional capability that the transmitter has the ability to pass the transmission privilege to other members of the group. Only one transmitter may be transmitting multicast messages at a given time. Closed groups in which restricted transmission privileges can be shared require that all group members know about the existence of the group and that there be a well-defined mechanism for transferring the transmission privilege among members of the group.

4.2.1.10 Closed Group, Unrestricted Transmission source

4.2.1.10.1 A closed group with unrestricted transmission defines the case in which only members of the closed group are permitted to originate multicast messages destined for the group. All group members, however, may simultaneously send messages to all other members. There is no designated source, but there may still be a need for a token passing mechanism to pass the transmission privilege among group members. All group members are peers, with the data sent by any one member going to all other members. If a closed group-based multicast transmission service were used to support teleconferencing, restricted transmission would be analogous to the case in which a conference participant is required to get the floor before transmitting. Unrestricted transmission would be analogous to the case in which conference participants could interrupt each other and transmit at the same time.

4.2.1.11 One-way, Two-way, N-way Transmission

4.2.1.11.1 The unrestricted transmission source scheme is also describable as an N-way transmission scheme. The fixed and single transmission source schemes can be further classified as either One-way or Two-way schemes, depending upon whether the group members receiving the multicast transmission from the source are permitted to send responses to the source. One-way describes a scheme in which group members are permitted only to receive transmissions from the source; Two-way describes a scheme in which group members may both receive transmission from the source and send data to the source.

4.2.2 **Group Size**

4.2.2.1 Certain applications may be designed for small groups of 1-5 recipients while others may be designed for medium-sized or very large multicast groups with the number of recipients in the thousands. The size of the group for which an application is designed may affect the choice of acknowledgement scheme, type of group membership, the ability to tolerate group member movement, and other multicasting service issues. An application designed for small groups may be more flexible regarding the variety of services that could be used to support it.

4.2.2.2 Stationary/Movable/Mobile Group Members

4.2.2.2.1 Multicast groups can consist of members that are stationary and remain attached to the network at the same subnetwork point of attachment, or they may have group members which can leave one point of attachment and reconnect to the network at a different point of attachment. In addition, they may have members who are in continuous motion who ant to retain their ability to send and receive multicast transmissions from other group members while they are in motion. Members that remain in one permanent location are called stationary members, members that can detach from the network and reattach at a different address location are called movable, and those which retain connectivity while in motion are known as mobile. The dynamic addressing required to send data to groups in which members are mobile and/or movable is necessarily more complicated than the addressing required to address. Most of this addressing complication, however, is independent of whether or not the data being transmitted is unicast or multicast, and is largely a separate problem to be solved.

4.2.3 Quality of Service

4.2.3.1 Unconfirmed Transmission

4.2.3.1.1 If multicast transmission is unconfirmed, the sending application will not receive any acknowledgment from either the service provider or any of the recipient peer application entities that a transmitted message was in fact received correctly by all intended recipients. This is the simplest case, providing transmission simplicity, efficiency, and flexibility at the expense of reliability and control. The sender need not know anything about the identities or even about the number of entities in the destination group. The sender merely sends one protocol data unit (PDU) to a group, and need not maintain any state information regarding which of the destination entities did or did not receive it, and which might require retransmission. Failure to maintain state information, however, also means that the sender has no way of keeping track of the number of active recipients in the destination group even when this number drops down to zero and there are no more active recipients. The sender has no way of detecting whether or not any recipient did in fact receive a multicast transmission, so he has no way of assuring which recipients receive what or of synchronizing communications among the group.

4.2.3.2 Negatively Acknowledged Transmission

4.2.3.2.1 Delivery of multicast PDUs may remain unacknowledged by any explicit means, yet be acknowledge implicitly by a transmission protocol that has been designed to provide reliability by ensuring that all transmission protocol that has been designed to provide reliability by ensuring that all transmission errors will be detected and, if possible, recovered from. The protocol reports all transmission errors to the message originator so that all lost data can be retransmitted. The absence of transmission error notification in such a multicast transmission may serve as implicit acknowledgment that all recipients received the transmitted multicast message without error. The absence of such negative acknowledgments, however, is not a guarantee of correct receipt by all recipients because those recipient nodes that have gone down or are otherwise malfunctioning do not have any way of transmitting a negative acknowledgment to indicate that they did not correctly receive the data. Hence, the negative acknowledgment service is not a fully reliable one.

4.2.3.3 Confirmed Transmission

4.2.3.3.1 Increased reliability and control may be obtained using the practice of having the receipt of multicast messages acknowledged. In such a case, the originating application or service provider receives and maintains state information regarding the receipt of such acknowledgments to determine whether to retransmit protocol data units, if necessary. Confirmation may come either from the recipient application entities themselves, or from the service provider serving them.

4.2.3.4 Ordered Delivery

4.2.3.4.1 If an ordered delivery service is provided, then all data transmitted via multicast is received by all recipient entities in the same order in which it was transmitted.

4.2.3.5 Degrees of Reliability

4.2.3.5.1 Reliability generally refers to a delivery service that guarantees that data which is transmitted via multicast is received error free and in the same order in which it was transmitted. There are several different degrees of reliability that can be attained, depending upon how many acknowledgment are required to be received. The confirmation scheme may involve the multicast originator receiving acknowledgments confirming receipt of the multicast transmission by all, some minimum number, some specific subset, or only one recipient(s). For example, the multicast transmission may be acknowledged by each of the remote service providers or each of the remote application entities upon receipt of the transmission at each of the multicast destination end systems. If generated by the destination application entities, each of the acknowledgments may be delivered to the multicast originator as confirmation. Depending upon the policy of the originating application, acknowledgments may be required to be received form all, some specific number, some specific subset, or only one of the destination entities. The originating application would be responsible for keeping track of the acknowledgments received in order to maintain state information regarding the number of active members, the status of the group, and the status regarding which members of the group responded to the transmission. If the acknowledgment are generated by the destination service providers, the duties of receiving and maintaining state information regarding such group and transmission status can be the responsibility of the service provider. In fact, the service provider can keep track of acknowledgments without having to pass them all to the originating application. Again, confirmation may be required in varying degrees, depending on the application being served. The application would need to interact with multicast group management facilities to define the specific confirmation policy. Confirmation information in the form of acknowledgments received enables the originator to determine whether to retransmit PDUs, if necessary, and to determine when the group no longer contains any active members.

4.2.3.6 Recovery from Lost Data and Other Transmission Errors

4.2.3.6.1 Given a confirmation service which enables the originator to detect the failure of a recipient to receive a multicast transmission, a multicasting service may or may not have the ability to recover form such errors. A reliable multicasting service is one in which transmission errors such as down connections and lost, reordered, or corrupted data can only be detected, but also recovered from via retransmission. Such a service is said to be fully reliable.

4.2.3.7 Synchronized/Unsynchronized

4.2.3.7.1 In the synchronized multicast, the recipient service providers withhold delivery of the received message to the recipient peer entities until a notification is received from the sender. Depending on the application, this notification could come from either the sending application process of the sending service provider. Such notification enables the sender to control, to some extent, whether all recipient peer entities receive and process the same information. A multicast service which does not have such a capability is said to provide only unsynchronized transmission.

4.2.4 Additional Quality of Service Factors

4.2.4.1 In addition to quality of service and other requirements that are specific to multicasting, an application may have other, additional requirements that are not specific to multicasting but that affect the appropriateness of each variety of multicasting service for use in supporting that application. Examples of such separate but related requirements include requirements for message confidentiality or other forms of data security, for the ability to work over both intrinsically multicast and non-multicast media, and for achieving a certain data throughput. Other capabilities and restrictions relating to the particular application being used and/pr the context in which it is used may also affect the appropriateness of each variety of multicasting service for use in supporting that application. The environment in which the application is functioning may affect how group addresses may be set up, how routing information may be distributed, what upper layer facilities may be required, or other factors. Such additional qualities of service affecting the choice of what selection of multicasting services are best suited to support particular applications are discussed in this subsection.

4.2.5 Data Type

4.2.5.1 An application may need to multicast many different kinds of data. For example, the data may consist of database updates, non-real time messages, voice, video, crucial tactical battle information, time-sensitive imagery, files, etc. The type of data to be multicast by an application will certainly affect the choice of multicasting services that should best be used to support it. Time-critical information will require highly efficient transmission, while non-real time information such as file transfer, database updates, and mail will tend to require reliable transmission.

4.2.6 Message Size

4.2.6.1 Message size, or the number of bytes in a typical message multicast by an application, may also affect the appropriateness of various types of underlying varieties of multicasting service for supporting that application. The reliability , delay, and other requirements of the application being equal, a multicasting service that will be transmitting predominantly short messages is probably best designed differently, in terms of error recover schemes and other quality of service factor, than a multicasting service that will be transmitting messages several megabytes long.

4.2.7 Frequency

4.2.7.1 Frequency refers to the number of messages per unit of time that the application typically multicasts. One application may multicast 100 messages per minute while another multicasts only one message per day. The frequency with which the application transmits multicast data will affect the traffic load of the underlying network and, therefore, should be a determining factor, along with the application's other requirements, in deciding what variety of multicasting service with which to support the application.

4.2.8 Throughput

4.2.8.1 Throughput refers to the number of bytes of user data that can be transferred through the network per second, as measured over some time interval. An application's throughput requirements will certainly be a determining factor in the type of multicasting service that would best support the application. Throughput is at the opposite end of the spectrum from reliability. As throughput requirements increase, the ability to provide reliable transmission decreases. Stringent throughput requirements will certainly influence the type of acknowledgment scheme, if any, that should be employed as part of the multicasting service.

4.2.9 Delay

4.2.9.1 Delay typically refers to the time between a message being sent by an originator and its being received by a recipient across the network. In the context of multicast transmission, however, the concept of delay becomes multi-faceted because there are multiple recipients for any given multicast message. If some recipients are significantly more distant from the source than others, both geographically and in terms of the number and type of different networks that must be traversed to reach the recipients, the delay experienced by the distant recipients may be significantly larger than the delay experienced by closer ones. An application may have a requirement that all, or only that some minimum subset, of the recipients receive messages within a certain delay time. As an application's delay requirements become more stringent and comprehensive, the underlying transmission protocol that serves the application is forced to limit the extend of the reliability, security, and other services it provides. Additionally, an application may have a requirement that the different delays experienced by recipients of a given transmission do not vary beyond a certain amount. In this case, especially if a group is spread out over a wide geographic area and across different types of networks, it may be difficult or impossible to guarantee that a given transmission will reach all recipients within a certain amount of time. In such a case in which very small variance in dleay is required but difficult to accommodate, it may be beneficial for the network to provide a synchronization service. A synchronization service would enable the recipient service providers to withhold delivery of a received message to the recipient peer entities until a notification is received from the sender, thereby decreasing delay variance at the expense of decreasing message throughput.

4.2.10 Network Type

4.2.10.1 Some transmission media, such as radio and local area network technology, inherently possess the ability to support communications to multiple destinations simultaneously. Others, such as point-to-point networks, do not. Applications may be required to work over such intrinsically multicast media or to work over intrinsically non-multicast media, or over an internetwork consisting of a concatenation of both types of subnetworks. The type of underlying media assumed may affect the appropriateness of a given variety of multicasting service for use in support of the application being considered. An application that must run over media that is not intrinsically multicast may be better serviced by a multicasting service that, for example, is based on a static membership policy with a fixed transmission service.

4.2.11 Addressing

4.2.11.1 For the most part, when an application initiates data transfer to a multicast group, the way in which the group is addressed should be transparent to the application because such addressing should be the responsibility of a network service underneath the application. If an application does have a specific requirement to be able to address the multicast group either by a group address or by an explicit list, or both, then this requirement should be taken into account when determining the variety of multicasting service to implement. If an application uses a group address, then the underlying multicasting service must include group management mechanisms for translating that group address into an explicit list of destination entities and their addresses.

4.2.12 Security

4.2.12.1 In order to take advantage of the benefits of increased efficiency and timeliness which multicasting provides, it is necessary that a single message be sent out to multiple recipients. Hence, all recipients receive an identical message from the source. If this message is to be protected cryptographically, while it is in transit over the intervening internetwork, then each of the recipients must posses the cryptographic key needed to decipher the message into its plaintext form. Hence, message confidentiality in such a situation requires that all members of a group be trusted to the same extent. An application requiring such message confidentiality may be best served by a multicasting service that, for example, is based on closed, static, and determinate membership in order to simplify the tasks of cryptographic key management.

4.2.13 Environment

4.2.13.1 The environment in which the application will typically operate may have special constraints or characteristics that also affect what multicasting services should best support it. "Environment" is a vague term that is meant to refer to situations such as those in which a typical multicast group encompasses geographically dispersed systems, only systems on a single subnetwork, systems dispersed over disparate internetworked subnetworks, systems operating in a tactical environment in which bandwidth is limited and communications unreliable, etc. The application's typical operating environment will determine the application's requirements, thereby determining in part what variety of multicasting service is best used to support it.

4.2.14 Application Conditions and Requirements

4.2.14.1 Applications that can best be enhanced by underlying multicasting data communications services are very often applications which are time and/or bandwidth critical. Multicasting provides a means by which an identical message or block of data may be transmitted only once from the source, yet reach the multiple destinations for which it is intended. Such a service provides efficiency by minimizing bandwidth, processing power and/or the transmission delay. The timeliness, bandwidth and processing power efficiency provided by multicasting is not without some tradeoff, however, for as efficiency increases, reliability tends to decrease. Similarly, as measures to increase reliability are introduced into the multicasting protocol, efficiency (time, processing power, and bandwidth savings) tends to decrease. The correct balance between efficiency and reliability must be evaluated for each application to be supported by multicasting services to determine the specific variety of multicasting to be used. Ultimately, our goal is to distill some core sets of multicasting services, each of which is beneficial to a specific category of applications, and earmark these for eventual incorporation into standards.

5 Status of Multicast Standards and Activities

5.1 The following is a list of the current OSI Multicast standards.

5.1.1 Amendment 1: 1995 to ISO/IEC 8473-1: 1994 provides multicast extensions to CLNP. This text is mostly editorial in nature so that ISO/IEC 8473 recognizes that multicast forwarding exists and is distinct from single cast forwarding.

5.1.2 ISO/DIS 9542 is in the progress of replacing ISO 9542: 1988 (the ES-IS protocol). This draft international standard incorporates both amendment 1 (NSAP Address Assignment) and Amendment 2 (Multicast Extensions).

5.1.3 Amendment 1:1996 to ISO/IEC 8602:1995 (Addition of connectionless-mode multicast capability) extends the use of the connectionless network layer multicast service to the transport layer.

5.1.4 ISO/IEC 8348:1996 (Network Service Definition) incorporates Amendment 5 (addition of Group NSAP Address).

5.2 ATNP WG 2 Proposal

5.2.1 A proposal was received in WG 2 for a modification to ES-IS for the purpose of adding multicast capabilities. The proposal included ATN specific enhancements to the ES-IS protocol for multicast functionality.

6 Analysis of Proposals

6.1 Analysis of Differences

6.1.1 The multicast amendment to ES-IS has been specified as a natural extension to ISO/IEC 9542, while the ATN specific approach was specified bearing in mind ATN requirements and this leads to the different styles and advantages and disadvantages of each approach:

a) ATN User Requirements tend to place great importance on knowing when communication has been lost. With the standard ES-IS approach, loss of membership of a multicast group can only be determined through non-receipt of data. When multicast communications are used for continuous or regularly data transfer (e.g. radar data distribution) this is not an issue. However, when it is used for irregular communications (e.g. DFIS type messages to a group of aircraft), then there is no easy way to detect loss of group membership.
b) Multi-homed End Systems are expected to be the norm in the ATN. The ATN specific approach was deliberately designed to avoid a multi-homed End System receiving multiple copies of the same multicast NPDU, and achieves this through the directed registration mechanism. Furthermore, if registration

is not possible through one subnetwork and router, the reporting of loss of membership ensures that the End System can try to gain access to a multicast group through an alternative subnetwork/router.

c) Under multicast ES-IS, if an End System sends out its End System Group Hello (ESGH) on multiple subnetworks then it can expect to see multiple copies of each NPDU. Indeed, there is no easy way for the Routers to recognise such a multiple homed ES and suppress multiple copies, because the ESGH does not include a System Level Identification (e.g. an NET), the only possibility is to infer such information from ESH PDUs - however, consistency between the NSAP Addresses reported by ESH PDUs on different subnetworks is not guaranteed.

6.1.2 If multicast ES-IS is used unmodified for the ATN, multi-homed ATN End Systems will have to broadcast their ESGH PDUs on all attached subnetworks in order to ensure high availability. However, in turn, they will typically receive multiple copies of the same NPDU sent to the multicast groups that system listens to. The ES will not be told when total loss of service occurs.

6.1.3 Multicast ES-IS appears to meet a requirement for best efforts general multicasting where any End System may send an NPDU to a given Group NSAP Address. However, the ATN specific approach specified a more constrained model of multicast communications (only one source End System for a given Group NSAP Address) which met a certain set of ATN user requirements (high availability without frequent NPDU duplication, user informed of loss of access to a given multicast group, etc.). and which were perceived as implementable. Multicast ES-IS may need to be extended to include features present in the ATN specific approach if it is to be used for ATN multicast. As it is desirable to keep as close as possible to ISO standards, such enhancements need to be identified and possibly forwarded to ISO as defect reports.

6.2 Enhancements Needed for Multicast ES-IS use in the ATN

6.2.1 Mulicast ES-IS enhancements are needed to avoid NPDU duplication with multi-homed End Systems and to report multicast group membership status when the multicast model is restricted to a single source for a given Group NSAP Address. These enhancements appear to be:

7. The ESGH PDU is extended to include an NET that uniquely identifies the source ES. This NET is included in every ESGH PDU sent by the ES and may be used by Routers (and by later extensions to ISO/IEC 10589) to avoid sending multiple copies of the same NPDU to an ES.
8. The concept of the “Active Multicast IS” and the Multicast Address Mapping (MAM) PDU does not appear to be appropriate to this model of multicast communications. Instead, a new PDU needs to be introduced which allows ISs

to respond to ESGH PDUs and report the multicast Group Addresses for which they are part of the multicast distribution tree, and the SNPA Addresses to which they will send them. This is similar to the Register Acknowledge PDU of the ATN specific approach.

6.2.2 The really important enhancement is the second one as will extend multicast ES-IS to operate in an ATN compatible manner. Once an ES has determined which Routers can satisfy its request to join a multicast group, it can then cease broadcasting its ESGH and send it unicast to that Router (thus hopefully avoiding duplicates). Should that Router ever stop sending out its PDU announcing that it supports that multicast group, then End System can start broadcasting the ESGB on all subnets in order to locate an alternative.

7 Future plans

7.1 Further investigation is needed to decide the addition of multicast to the ATN protocols.