ETSI TECHNICAL REPORT

ETR 260

May 1997

Source: ETSI TC-NA

Reference: DTR/NA-021206

ICS: 33.020

Key words: Addressing, ATM, CBDS, MAN

Network Aspects (NA); Connectionless Broadband Data Service (CBDS); Addressing principles and related aspects for the CBDS

ETSI

European Telecommunications Standards Institute

ETSI Secretariat

Postal address: F-06921 Sophia Antipolis CEDEX - FRANCE
Office address: 650 Route des Lucioles - Sophia Antipolis - Valbonne - FRANCE
X.400: c=fr, a=atlas, p=etsi, s=secretariat - **Internet:** secretariat@etsi.fr

Tel.: +33 4 92 94 42 00 - Fax: +33 4 93 65 47 16

Contents

Foreword

This ETSI Technical Report (ETR) has been produced by the Networks Aspects (NA) Technical Committee of the European Telecommunications Standards Institute (ETSI).

The purpose of this ETR is to provide guidance on addressing principles and related aspects for the Connectionless Broadband Data Service (CBDS).

ETRs are informative documents resulting from ETSI studies which are not appropriate for European Telecommunication Standard (ETS) or Interim European Telecommunication Standard (I-ETS) status. An ETR may be used to publish material which is either of an informative nature, relating to the use or the application of ETSs or I-ETSs, or which is immature and not yet suitable for formal adoption as an ETS or an I-ETS.

Blank page

1 Scope

This ETSI Technical Report (ETR) focuses primarily on the addressing principles required for CBDS although some additional closely related aspects are also included to assist clarity and ensure coherence.

Clauses 1, 2 and 3 of this ETR gives an introduction, references and definitions, clause 4 an overview of Metropolitan Area Networks (MANs) as described in ETS 300 275 [5], Asynchronous Transfer Mode (ATM) networks and CBDS, and clause 5 deals specifically with addressing principles for the CBDS. Group addressing is dealt within clause 6.

2 References

This ETR incorporates by dated and undated reference, provisions from other publications. These references are cited at the appropriate places in the text and the publications are listed hereafter. For dated references, subsequent amendments to or revisions of any of these publications apply to this ETR only when incorporated in it by amendment or revision. For undated references the latest edition of the publication referred to applies.

[1] ETS 300 211: "Network Aspects (NA); Metropolitan Area Network (MAN); Principles and architecture".

[2] ETS 300 212: "Network Aspects (NA); Metropolitan Area Network (MAN); Media access control layer and physical layer specification".

[3] ETS 300 217-1: "Network Aspects (NA); Connectionless Broadband Data Service (CBDS); Part 1: Overview".

[4] ETS 300 217-4: "Network Aspects (NA); Connectionless Broadband Data Service (CBDS); Part 4: Address screening supplementary service".

[5] ETS 300 275: "Network Aspects (NA); Metropolitan Area Network (MAN); Interconnection of MANs".

[6] ETR 112: "Broadband Integrated Services Digital Network (B-ISDN); B-ISDN Principles".

[7] ITU-T Recommendation I.150: "B-ISDN asynchronous transfer mode functional characteristics".

[8] ITU-T Recommendation I.361: "B-ISDN ATM layer specification".

[9] CCITT Recommendation E.164: "Numbering plan for the ISDN era".

[10] ETR 122: "Network Aspects (NA); Connectionless Broadband Data Service (CBDS); CBDS over Asynchronous Transfer Mode (ATM)".

[11] ETR 073: "Broadband Integrated Services Digital Network (B-ISDN); Evolution towards B-ISDN".

3 Definitions and abbreviations

3.1 Definitions

For the purposes of this ETR, the following definitions apply:

CCITT Recommendation E.164 [9]: CCITT Recommendation E.164 [9] provides a coherent and flexible numbering plan that can be applied to all public telecommunications services in the ISDN era. This will enable easy identification of all ISDN and PSTN network interfaces and the interconnection of services that are compatible in other respects.

ConnectionLess Server (CLS): It is the network element containing the ConnectionLess Server (CLS) functional group, see ETR 122 [10], subclause 4.3.

Inter Network operator Interface (INI): Interface between different operators networks which are directly connected

MAN Switching System (MSS): A collection of functions that provides high-speed switching in the public network. It can be implemented through distributed or centralized switching, see ETS 300 211 [1].

numbering scheme modelled on CCITT Recommendation E.164 [9]: A numbering scheme that aligns with the numbering structure and criteria as specified in CCITT Recommendation E.164 [9], unless specifically amended within the context of this document.

Protocol Data Unit (PDU): The PDU is a block of data which consists of all information (user data, addressing information and service parameters) related to a single, self contained service instance.

Service Switching System (SSS): A service switching system is either an MSS or a CLS.

Spanning Tree (ST): Any bi-directional communication network can be modelled as an un-directed graph where the nodes are the service switching systems and the arcs the bi-directional links between them. A Spanning Tree (ST) over all the nodes is a subset of the arcs such that the resulting graph contains no loops.

time "T": This refers to the specified time after which all PSTN and ISDN networks can use the full capability of CCITT Recommendation E.164 [9].

User MAN Interface (UMI): Interface between the CN and the Access Facility 1 (AF1), see ETS 300 211 [1].

User Specific Interface (USI): Interface based on a user specific protocol, see ETS 300 211 [1].

3.2 Abbreviations

For the purposes of this ETR, the following abbreviations apply:

AF1	Access Facility 1
ATM	Asynchronous Transfer Mode
B-ISDN	Broadband Integrated Services Digital Network
CBDS	Connectionless Broadband Data Service
CC	Country Code
CNM	Customer Network Management
CNMA	Customer Network Management Agent
CLAI	ConnectionLess Address Identifier
CLNAP	ConnectionLess Network Access Protocol
CLS	ConnectionLess Server
DN	Data Network
DQDB	Distributed Queue Dual Bus
GA	Group Address
GAA	Group Address Agent
GAP	Group Address Packet
IA	Internal Address
IEEE	Institute of Electrical and Electronic Engineers
INI	Inter Network operator Interface
ISDN	Integrated Services Digital Network
LAN	Local Area Network
MAN	Metropolitan Area Network
MSS	MAN Switching System
NDC	National Destination Code
NGAA	Nested Group Address Agent
NGA	Nested Group Address
NMS	Network Management Station
NNI	Network Network Interface
OS	Operation System
PDU	Protocol Data Unit
NSN	National Significant Number

PSTN	Public Switched Telephone Network
SMDS	Switched Multimegabit Data System
SN	Subscriber Number
SSS	Service Switching System
ST	Spanning Tree
TE	Terminal Equipment
UMI	User MAN Interface
USI	User Specific Interface

4 CBDS and the supporting technology

4.1 Basic principles of Metropolitan Area Networks (MANs)

Metropolitan Area Networks (MANs) are based on the use of a shared access broadband medium and are used in a multi-user environment for high speed data transfer which is oriented to public domain network applications. Typically a MAN may link private networks in the form of Local Area Networks (LANs) to public networks such as narrowband (and in the future broadband) ISDNs. MANs provide the integrated support of narrowband and broadband services which could be data, voice or video, although early applications are targeted at data services. A MAN typically covers an urban or metropolitan area 50 km wide although several MANs could be linked together. Applications are normally targeted towards the large business sector of the Telecommunications market providing integrated support of a range of service offerings (data, voice, video) as supported by bearer services e.g. connectionless, connection-oriented isochronous and connection-oriented non isochronous services as described in ETS 300 211 [1].

4.2 Basic Principles of ATM

ATM is the transfer mode for implementing B-ISDN. The basic principles of ATM are defined in ETR 112 [6], ITU-T Recommendation I.150 [7] and ITU-T Recommendation I.361 [8].

The provision of CBDS service in B-ISDN is realized by means of ATM switched capabilities and Connectionless Service Functions as described in ETR 122 [10], subclause 4.2.

4.3 Evolution of MANs

Man technology based on the Distributed Queue Dual Bus (DQDB) protocol will provide the initial CBDS offering for those users who choose to take the service in the first instance. This capability will provide LAN-to-LAN and LAN-to-HOST interconnection by dedicated links. The second stage will see Asynchronous Transfer Mode (ATM) interconnection of MAN Switching Systems (MSSs) over an ATM interface. Connection will be established on a semi-permanent basis through ATM digital cross connect switches. The third evolutionary phase could occur with the widespread introduction of B-ISDN services, when MANs could serve as gathering networks for a fully switchable ATM based network.

4.4 Evolution of ATM networks

The evolution of ATM networks taking into account relevant developments and trends in technical and non-technical sectors is described in ETR 073 [11].

4.5 Introduction to CBDS

A CBDS can be provided by MANs and Asynchronous Transfer Mode (ATM) based networks to extend LAN like performance (see note*)* across a metropolitan or wide area. Access could be at 2, 34, 140 or 155 Mbits. This enables Protocol Data Units (PDUs) to be transferred across the network without the establishment of end to end connection.

NOTE: However, the inherent delay of long distance communications could degrade end to end performance as seen by the user if end to end communication protocols are not adequate.

5 Addressing principles

5.1 Addressing requirements

Each PDU sent by an end user has a source address and a destination address which identifies the sender and intended recipient respectively. These requirements form the basis of both individual and group addresses. Addresses are modelled according to CCITT Recommendation E.164 [9].

5.2 Definition of addresses and address types

The following address types are utilized within CBDS:

5.2.1 Individual address

An individual address could be either a source or destination address identifying a user. It represents the address of a particular interface (see note) at the T-reference point. More than one number may be assigned to this interface at the "T" reference point (and could be used as an address).

NOTE: The interface through which the user accesses the service.

Source Address:

This is an individual address modelled on the CCITT Recommendation E.164 [9] numbering plan.

Destination Address:

This is an individual address modelled on the CCITT Recommendation E.164 [9] numbering plan.

5.2.2 Group Address (GA)

A group address is used as a destination address where a number of recipients are intended, each recipient being accessed through the use of the unique "group" identity. It therefore represents a set of individual addresses. Each individual address still has a CCITT Recommendation E.164 [9] type address but a 60 bit group address reference is used to determine multiple destinations. The intended recipients may be served by more than one network. In the transport of a PDU a Group Address (GA) can only be used a as destination address. A particular interface (see note above) at the "T" reference point can be identified by more than one group address administered by a single network. An interface at the "T" reference point is identified by a group address if one or more of the individual addresses assigned to this interface (see note above) at the "T" reference point is identified by the group address. The multiple destination transport for group addressing is specified in ETR 122 [10], subclause 4.1.1.

5.2.3 Nested Group Address (NGA)

A Nested Group Address (NGA) is related to a GA and identifies a subset of individual addresses pertaining to that GA. The NGA represents a set of individual addresses of a certain number of members of a GA, that are located inside a given network; this network can be different from the network where the Group Address Packet (GAP) originated, and from the network resolving the GA (the resolution function provides for a given GA or NGA, the list of all members addresses and/or NGA in the case of a GA).

An NGA is globally unique. A given NGA associated with a GA cannot be reused for another GA in order to allow independent evolution of the two GAs. A user should not put a NGA as a destination address in a ConnectionLess Network Access Protocol (CLNAP)-PDU, since a NGA is solely for use in PDUs exchanged between domains. It is an optional responsibility of each network domain to limit the effect of such a use of a NGA by a user, within their own domain. The use of NGAs can increase the efficiency of multicast transport by reducing the number of packets sent across networks.

5.3 Definition of a Group Address Agent (GAA) and Nested Group Address Agent (NGAA)

5.3.1 Group Address Agent (GAA)

The administrative aspects of any group address may be undertaken by a single GAA. Only one GAA will be responsible for administering any one group address. The GAA will be responsible for assigning, deleting, amending group addresses and for inclusion, addition and deletion of individual addresses to the group according to the instructions of the user/client or representative.

From the numbering point of view the GAA will assign a globally unique group address. The group address belongs to the GAA's address domain. The address domain generally is allocated in a national numbering scheme within the CCITT Recommendation E.164 [9] numbering plan.

The GAA represents a complete or partial group address resolution function. In the case of a partial address resolution by the GAA the complete resolution is accomplished with the support of other networks.

NOTE: There is a limit to the number of group addresses that a GAA can manage. Currently for the Switched Multimegabit Data Service (SMDS) it has been specified as 1 024 group addresses. The handling of group addresses above this number is for further study. In addition it is recommended that the maximum number of individual addresses and/or NGAs represented by a group address is 128.

5.3.2 Nested Group Address Agent (NGAA*)*

The NGA concept is applicable to encapsulating network; its applicability within non-encapsulating network requires further study.

The administrative aspects of a nested group address may be undertaken by a single NGAA. Only one NGAA will be responsible for administering a specific nested group address. The NGAA will be responsible for assigning, deleting, amending individual addresses associated with the nested group address in co-operation with the GAA which is responsible for the group, and according to the instructions of the user/client. The NGAA enables a stepwise and distributed group address resolution. It represents a group specific partial address resolution function in a network. This ensures the delivery of a group address packet to all members of the nested group address.

From the numbering point of view, the nested group address belongs to the address domain of the network to which the NGAA pertains and is globally unique. The NGAA represents a partial group address resolution function.

5.4 Address structure and format

5.4.1 Principles and structure of CCITT Recommendation E.164

The structure for CBDS addresses is modelled on CCITT Recommendation E.164 [9] which is shown in figure 1.

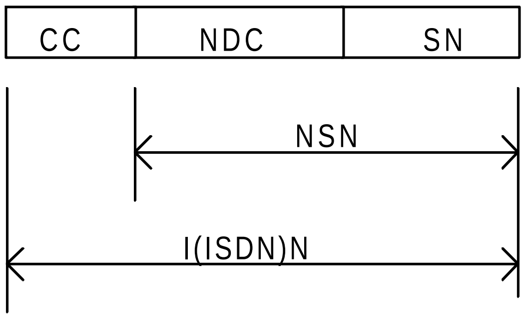

Where:

CC = Country Code;
NSN = National Significant Number;
NDC = National Destination Code;
SN = Subscriber Number;
I(ISDN)N = International ISDN Number.

Figure 1: Structure of CCITT Recommendation E.164 [9]

Within CCITT Recommendation E.164 [9] the international number length is variable with a maximum length of 15 digits after the 31st December 1996 (time "T"), with a 12 digit maximum applying until then. As the CBDS will not be switched within the PSTN the 12 digit restriction need not apply to this service and a full 15 digit capability is therefore available from the first day of this service. However, great care is needed when national variants of the addressing structure are implemented. The 60 bit (or 15 digit) finite length of the address subfield demands that all evolutionary aspects of a particular scheme are given due regard when the initial addressing requirements are introduced so that additional later requirements do not exceed this limitation.

5.4.2 Addressing format for CBDS

CBDS is one of several services provided by B-ISDN, and should follow a unique numbering plan. The CCITT Recommendation E.164 [9] number is allocated to the "T" reference point (as stated in ETS 300 217-1 [3], subclause 6.1). The same number is used in the MAC layer for MAN application to identify the MAC layer user entity, and in the CLNAP-PDU for ATM based network to identify the CLNAP user entities.

Destination and source address fields each have the same format. Address fields are 8 octets (64 bits) long, each address field having two subfields comprised of address type (4 bits) and address (60 bits). The address subfield is encoded with the binary coded decimal value of the number digits. CBDS address subfields are shown in figure 2.

Address Type	**Address**
4 bits	**60 bits (15 digit max.)**

Figure 2: CBDS address subfields

The address type subfield can be used to indicate either an individual or group address. The 60 significant bits available in the following address subfield provide the capability to carry a maximum number length of 15 digits which fully aligns with the post time "T" requirement of CCITT Recommendation E.164 [9] (see subclause 5.4.1).

NOTE: The address type of a CBDS address is sufficient to recognize a group address from an individual address. The address part of a CBDS address can in principle be identical for a group address and an individual address, however, a domain address space shall not be shared between group addresses and individual addresses in order to maintain the unambiguity of the CCITT Recommendation E.164 [9] numbering plan. However, an address shall exclusively be used as an individual or a group address within an address domain. To ensure that broadband services can evolve in a manner that will not be constrained by different addressing functionality being provided on a per service basis, e.g. the ability to recognize a group address from an individual address type, there should be no duplicate use of the same address.

For an individual address two possible forms have been identified. The first is full integration with the telephone/ISDN numbering i.e. a CBDS users number is not distinguishable from a telephone/ISDN users. The second form is the allocation of one or more specific geographic/non-geographical areas or number series, distinct from, and not utilized by telephone/ISDN users. The option selected remains a national matter.

A group address will not address an individual user, but will be used as a destination address to route to an entity capable of interpreting the group address or unique "group" identity, into the individual addresses of members or NGA addresses of that group.

5.4.3 Address type subfield

This is used to identify an address as either an individual or a group address. Values used to identify types of addresses are shown in figure 3.

Value	Address type
1100	Individual 60 bit publicly administered address
1110	Group 60 bit publicly administered address

NOTE 1: Only the values relating to CBDS are shown. The full list can be found in ETS 300 212 [2].

NOTE 2: An address type value to identify a nested group address is for further study; unless such a value is defined, the group 60 bit publicly administered address (1110) will be used.

Figure 3: Address type subfield values

5.4.4 Address subfield

The address subfield is encoded with the binary coded decimal value of the number digits which are used to route PDUs across Metropolitan Area Networks and/or ATM based networks. The address subfield is structured as shown in figure 4.

Country Code	National Significant Number

Figure 4: Address subfield

As detailed within CCITT Recommendation E.164 [9] the National Significant Number (NSN) is structured into two separate elements, the National Destination Code (NDC) and the Subscribers Number (SN). This structure used both in MAN and ATM based networks is shown in the example in figure 5.

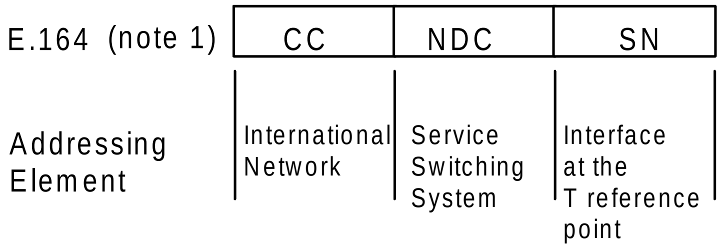

NOTE 1: E.164 = CCITT Recommendation E.164 [9]

NOTE 2: Whether a node can be identified by more than one CC+NDC combination is for further study.

Figure 5: Example of the possible structure of an address modelled on CCITT Recommendation E.164 [9]

In line with common practice when formulating numbering schemes the aim is to provide a structure which provides maximum flexibility and growth whilst minimizing overheads. To meet this criteria careful consideration is required by administrations when determining the National Significant Number (NSN).

Country Code element (CC):

The Country Code (CC) will be allocated in accordance with annex A of CCITT Recommendation E.164 [9] and has a maximum value of 3 digits.

National Destination Code element (NDC):

The NDC element is variable in length depending upon the requirements and may contain information related to the Service Switching System (SSS).

In line with CCITT Recommendation E.164 [9] the specific allocation and structure of digits used within the NDC is a national matter. However, it is recommended that each MSS/CLS should be identified within the first six digits of the NDC following the Country Code element (CC).

Subscriber Number element (SN):

The SN may vary in length but will identify the position of individual equipment within a sub-network.

5.5 Addressing requirements for both encapsulated and non-encapsulated data

For addressing purposes both data encapsulation or non-encapsulation is to be accommodated inside a network. With encapsulation an extra header is added to the PDU sent by a user as detailed in ETR 122 [10]. The extra header provides additional source and destination CBDS addresses for addressing and routeing purposes.

As a network option, within a single network operator domain PDUs whether addressed to a single user or to a group may be transported using either technique, i.e. encapsulation or non-encapsulation. For NNI applications between different network operators domains, PDUs crossing the NNI need to be encapsulated according to the technique defined in ETR 122 [10].

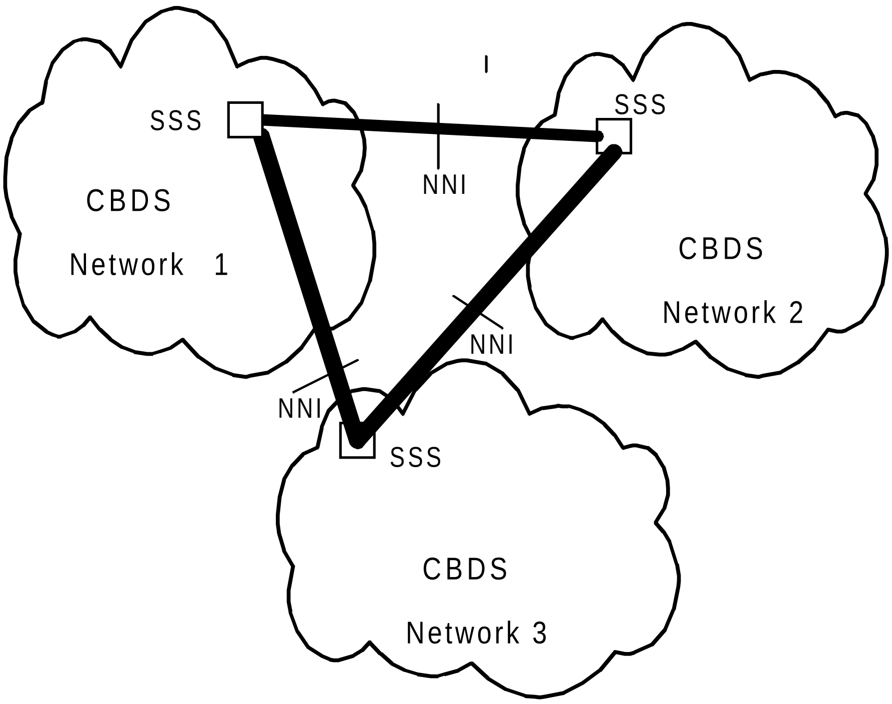

Figure 6: Example showing the possible interconnection of SSSs across an NNI where encapsulated group addressing needs to be used

Figure 6 shows three networks interconnected. All PDUs (both individual and group addressed) which cross an NNI need to use encapsulation.

5.6 Allocation of numbering resource

CCITT Recommendation E.164 [9] numbers provide the basis for global addressing capabilities for the ISDN. As such they exhibit certain qualities including;

- a hierarchical structure (CC NDC SN);
- a maximum length (15 digits following time "T");
- they are unambiguous.

The CCITT Recommendation E.164 [9] structured numbers for CBDS also need to comply with this criteria. Although there will be no direct connection with the PSTN/ISDN to meet the above criteria, numbers for this service should be allocated from within the national PSTN/ISDN numbering plan.

The address information carried in the address subfield shall always be the format shown in subclause 5.4.1, i.e. in the full international format (CC NDC SN) rather than the national format of NDC+SN. If more than one address format was used (requests for address information may well extend across international boundaries) then a method of determining type (or format) of address would be required but this is not recommended. It is also a requisite that PDUs will be transferred end to end unchanged, if an international format was not used the source address field would require modification to include the Country Code (CC) each time a PDU crossed the national boundary thereby infringing this requirement.

5.7 Network operator selection capabilities

Although CBDS will be initially introduced on a national level it will evolve into a pan-European/global network and as such network operator selection capabilities will be required to provide the required level of functionality as well as to align with the regulatory drive for competition. Network operator selection will also be required to enable a high degree of interconnectivity to be achieved as not all networks will be fully interconnected. Users will also want to control their service in a manner that allows them to take full account of the differing levels of service offered which may vary in terms of quality of service, performance and cost.

However, the need for further study of network operator selection on a worldwide level and not only as an issue related to a specific network or service is recognized. The study of network operator selection is currently not mature and on that basis it should be the subject of a separate ETSI document.

5.8 Address screening

Address screening is a supplementary service which enables the delivery of PDUs to be restricted. A full description of this service is given in ETS 300 217-4 [4].

6 Group addressing

6.1 General principles

Group addressing enables users to send a single data unit to multiple destinations, by identifying each required destination with a special destination address which is common to all required recipients. It is then not necessary to explicitly identify the individual address of each recipient in order to forward the packet until group address resolution has taken place. The group addressing feature provides a LAN like multicast capability:

- a group address is identified as such within the address type subfield in the destination address;
- a group address cannot be used as a source address;
- members of a group may be served by different networks;
- non members of a group address may send PDUs to that group;
- group addresses need to be unambiguous locally, nationally and internationally;
- group addresses need to be universally managed in a manner that conforms to CCITT Recommendation E.164 [9];
- group addressing provides a LAN like multicast capability over wide geographic areas;
- whenever group addressing is supported by a Spanning Tree (ST) on the network(s), care needs to be taken to ensure that the Spanning Tree (ST) will prevent loops which could overload the network(s) and lead to erroneous duplication(s) of PDUs.

6.2 Format

The format of a Group Address should follow the same format as an individual address and align with the CCITT Recommendation E.164 [9] structure (CC NDC SN).

6.3 Methodologies

Four alternative approaches are identified:

- a distributed database approach;
- a replicated database approach;
- a centralized database approach;
- a centralized approach with nested group addressing.

Each of these is now considered in more detail.

Compatibility of the different methods is for further study. However, the centralized database approach and the nested group addressing approach are compatible, see annex B.

NOTE: Detailed construction and compilation of routeing tables is network dependent and beyond the scope of this ETR.

6.3.1 Distributed database

6.3.1.1 Description

The resolution of group addressing using this approach requires each operator domain (i.e. a network which is either MAN or ATM based) to have its own data base which needs to hold all the information directly related to its own group members. Routing is achieved by utilizing a semipermanent Spanning Tree (ST) (see note 1) per address group which consists of the routes between the networks involved in the group address resolution. When a network receives a group addressed PDU, it resolves the group address as far as the users served by that network are concerned. In addition it sends the user provided group address packet (see note 2) to its neighbours in the Spanning Tree (ST). However, the data unit is not sent back to the originating user (if locally generated) or network. An example of the distributed database approach is depicted in figure 7. It indicates the users belonging to a group and the routeing information that has to be stored about the group in local (network) databases. A group address agent is still required in the distributed approach to administer the group address data at the internetwork level (e.g. to check the limitation on the maximum number of members of a group address is not exceeded). For routeing a group addressed packet from a non member user not located on the Spanning Tree (ST), a network may choose to route directly to the GAA or to the nearest network on the Spanning Tree (ST).

NOTE 1: For each group address, a related spanning tree has to be built over the nodes which have a knowledge of this group (the group address agents or nested group address agents) and the arc which connects them.

NOTE 2: A spanning tree arrangement does not require encapsulation of data, althought this need to be done between operators.

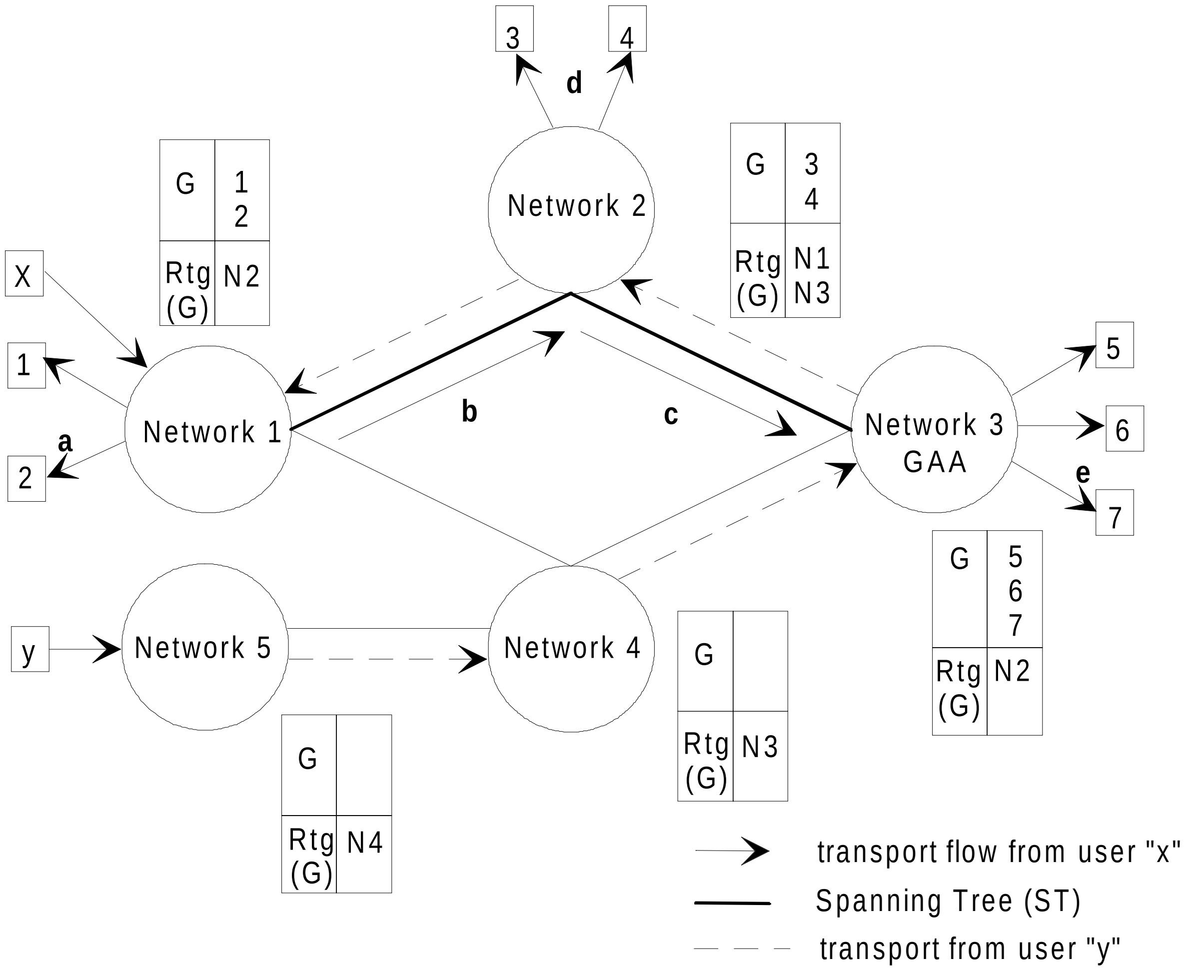

Tables show the information held by each of the distributed databases.

1 Members of the group served.

2 Routing information.

Figure 7: Distributed database approach

Five separate networks are shown, of which some users of networks 1, 2 and 3 are members of a particular group. Information about members of the group is distributed across each of the networks. When user "x" (on Network 1) wants to send a message to the members of that group the following action takes place. The group addressed PDU is passed across the network access and validated as a group address. The Network 1 routeing function identifies group members it directly serves (1 and 2) and resolves the group address into the individual addresses for each of those members in order to deliver the packet (a). The group addressed packet is then passed forward by its neighbours on the Spanning Tree (ST) except the one from which the packet was received (when it is the case) to Network 2 (b) where the same action is repeated (group address is resolved into individual addresses for members 3 and 4 and the packet delivered (d)). This process would be repeated for each hop across the Spanning Tree (ST) i.e. in the example also for Network 3 (c, packet (e)). In the configuration shown users "x" and "y" are not members of the group. Summarizing each group addressed packet is flooded to the networks concerned through a Spanning Tree (ST). The Spanning Tree (ST) allows optimization of the transport of multicast packets.

6.3.1.2 Adding a new member to the group

To add a new member to the group each of the distributed databases directly affected has to be updated to hold the relevant information. The action required will be determined by the Spanning Tree (ST) configuration employed. Examples are shown in figures 8a and 8b.

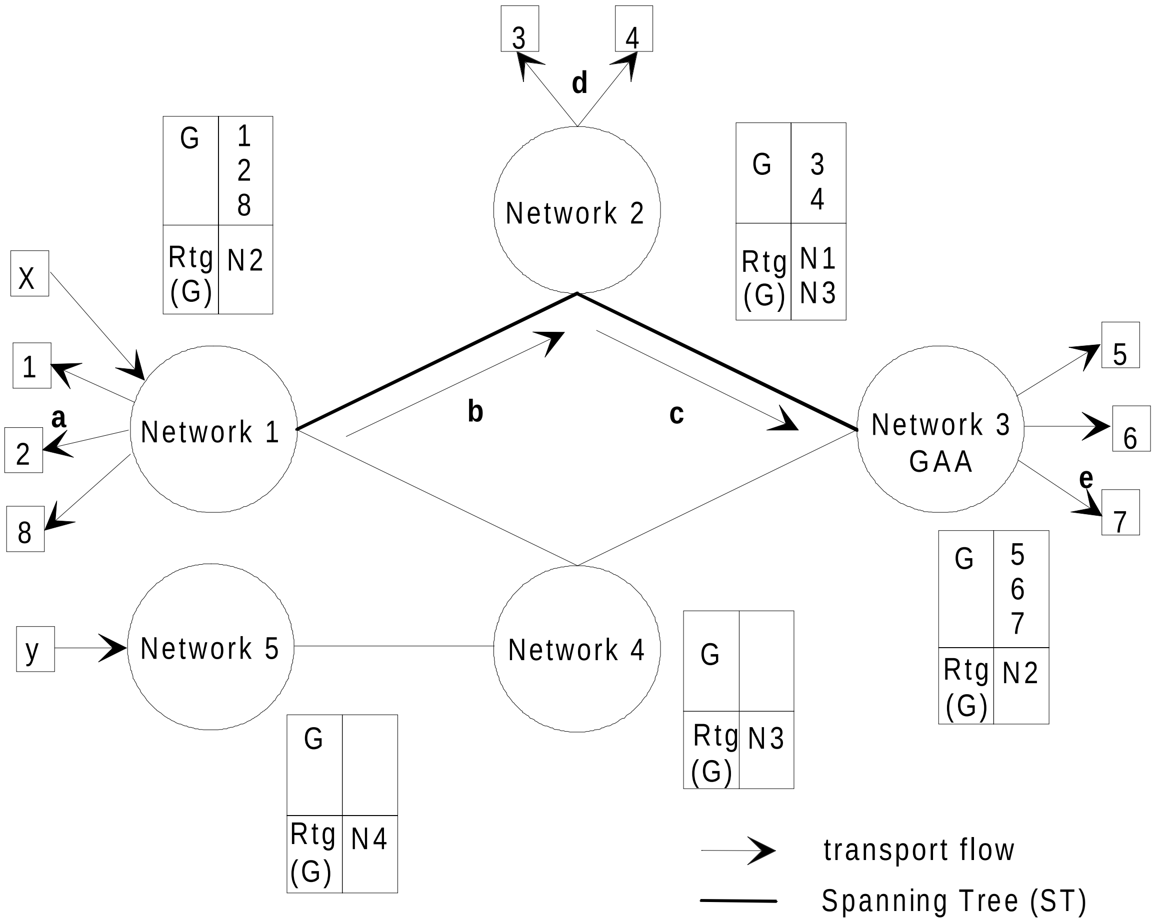

Figure 8a: A new group member to be added locally

For user 8 on Network 1 to be added to the group as a member, database information identifying user 8 as a member of the group and routing details will have to be programmed at Network 1. This is a local matter of Network 1 and does not affect the Spanning Tree (ST), i.e. the internetwork arrangements.

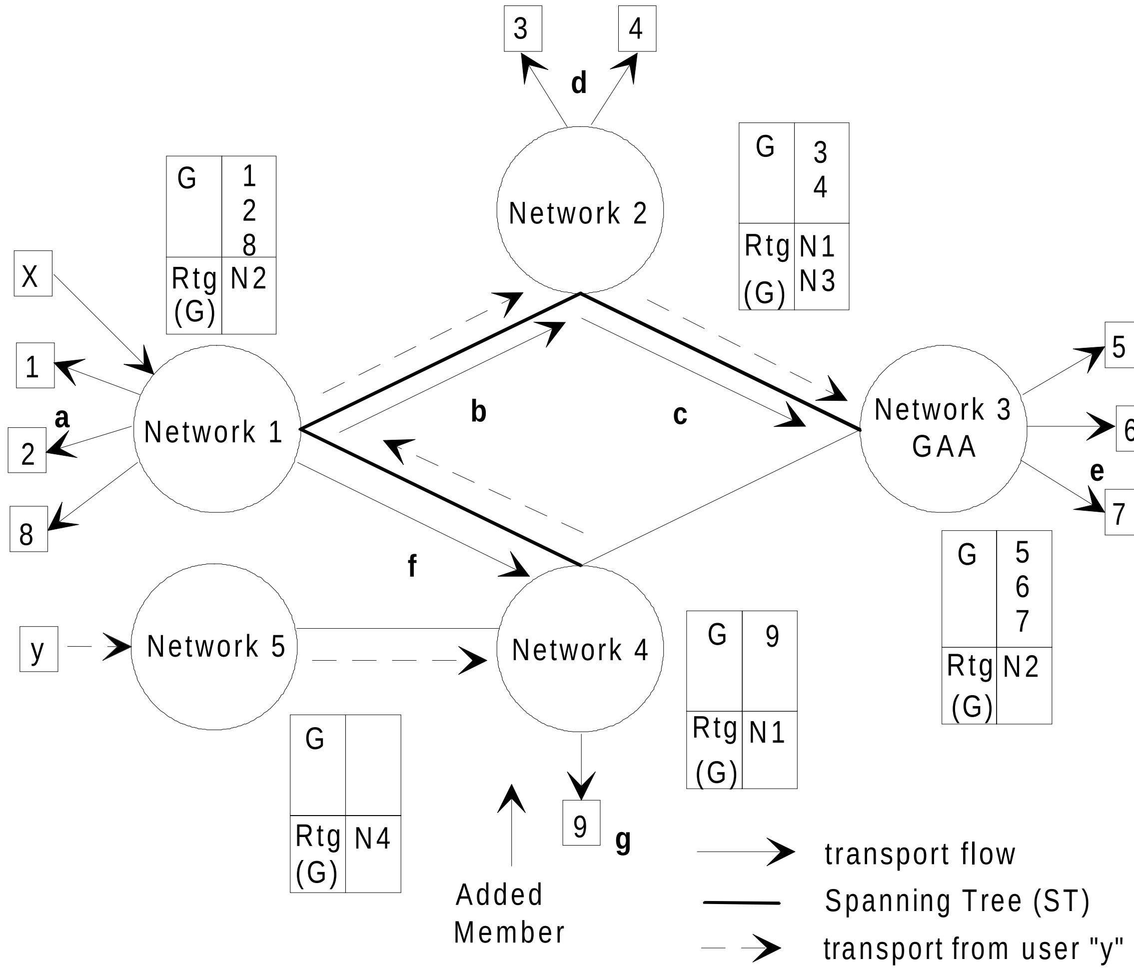

Figure 8b: A new group member at another network

For user 9 on Network 4 to be added to the group, database information identifying user 9 as a member of the group and routing details will have to be programmed at Network 4. The Spanning Tree (ST) routing shown also requires the routing information at Network 1 to be updated to include the hop to Network 4. This requires an additional internetwork arrangement to be established for the group. when user x from Network 1 wants to send a message to the members of the group, the following action takes place in addition to the actions described for figures 7 and 8. Network 4 is added to the Spanning Tree (ST), and the group address packet (f) is passed to Network 4 in order to deliver the packet (g) to member 9. Whether and when the Spanning Tree (ST) needs to be reconfigured is for further study.

6.3.2 Replicated database

This subclause is for further study.

6.3.2.1 Description

The resolution of group addressing using this approach requires each network domain to have its own database (it is not necessary to hold a database at each node) which needs to hold all the information relating to all Group Addresses (GAs) regardless of their location. For networks which do not directly have a member of the GA, it is sufficient to route the GAP to the GAA. Three cases can be foreseen A, B, and C:

A) a GAP is originated in a network that has a GA database. The GA is resolved in this network;

B) a GAP enters a network, and resolution did not take place along the path followed by this GAP. If this network has the GA database that resolution takes place in this network. This case requires routeing of the group addressed PDU to be achieved using a semipermanent Spanning Tree (ST) per address group which consists of the routes between networks involved in that group resolution;

C) the GAP reaches the GAA without being resolved, the GAA resolves the GA. Routeing of the group address PDU is achieved using a semipermanent Spanning Tree (ST) per address group which consists of the routes between the networks involved in the group address resolution.

An example of case A is depicted in figure 9.

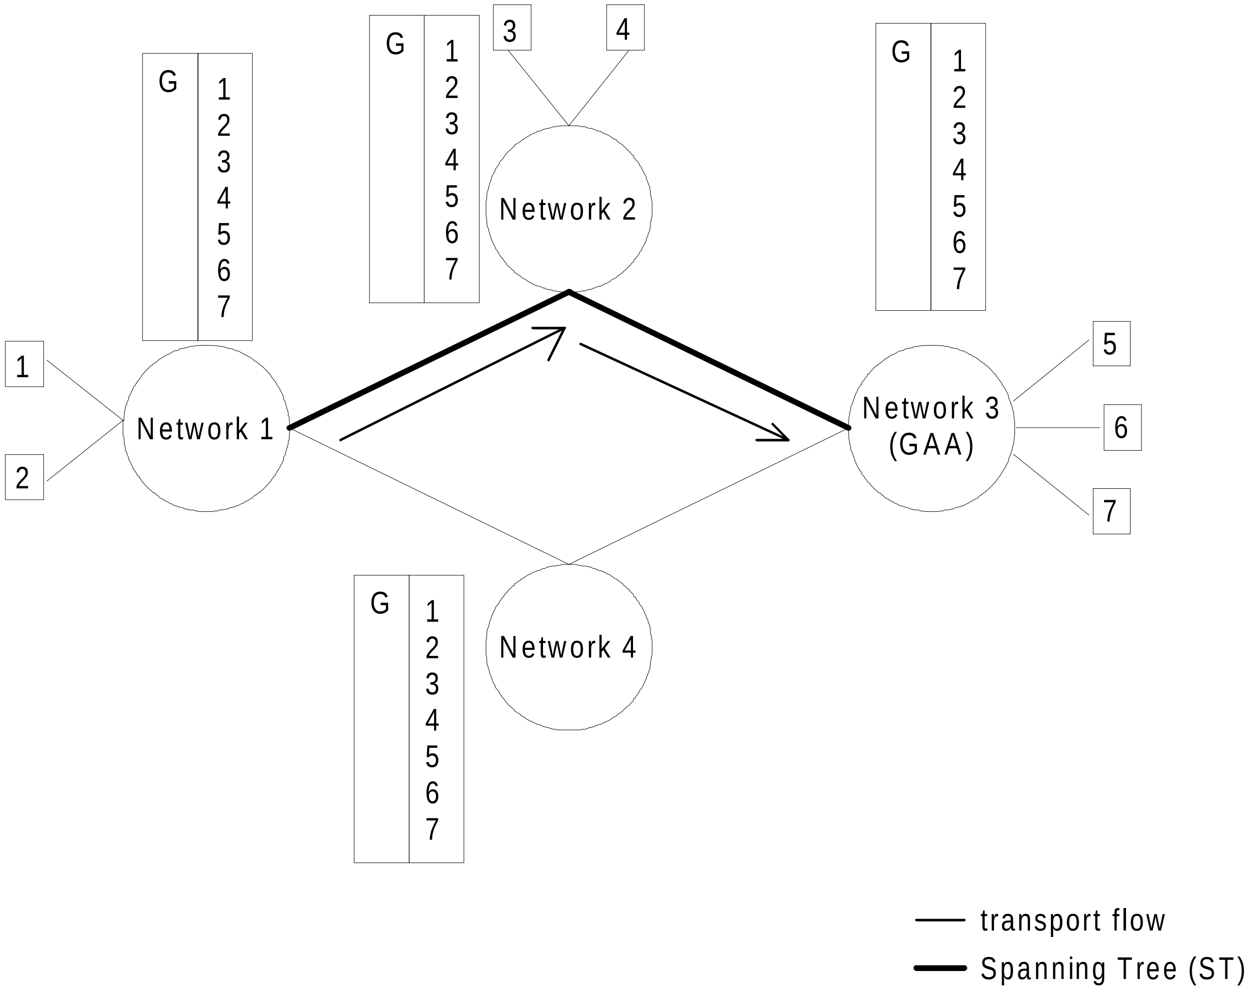

Figure 9: Replicated database approach - group address resolution functions

Four separate networks are shown, of which all the users of Networks 1, 2 and 3 are members of a particular group. When user 1 (on Network 1) wants to send a message to other members of that group the following action takes place. The group addressed PDU is passed across the network access and validated as a group address. The Network 1 database resolves the group address into individual addresses for distribution to the other networks. In summary each group addressed packet is resolved into individually addressed packets within the originating network and is then flooded through a Spanning Tree (ST) with the ingress network (or switching system) as the root.

6.3.2.2 Adding a new member to the group

To add a new member to the group each network domain has to update its database independently if it contains a member of that particular group. For example in figure 10, user 8 of Network domain 4 is added to an existing group containing members within Network domains 1, 2 and 3. Database information will need to be amended to reflect this change within each domain affected i.e. 1, 2, 3 and 4. Such actions will need to be strictly co-ordinated. Figure 10 shows Network 3 as the Group Address Agent (GAA). This function is required to control and co-ordinate changes as well as to ensure the limitations placed on the size of a group are not exceeded.

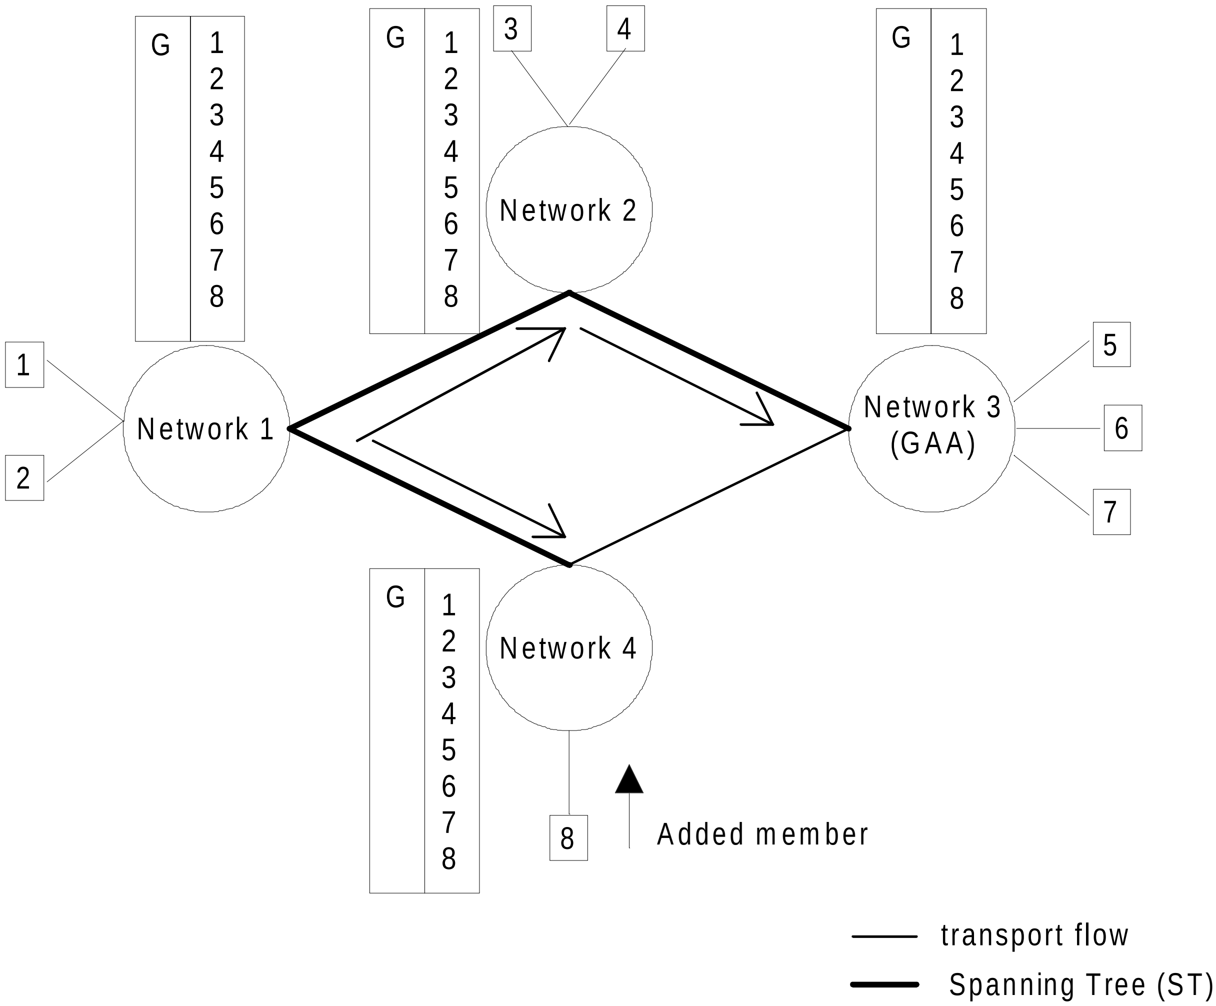

Figure 10: Replicated database approach - adding a new member

6.3.3 Centralized database approach

6.3.3.1 Description

A centralized database approach requires the use of a GAA. The GAA is responsible for administering, (assigning, amending, deleting) both group and individual addresses for the group. Only one GAA should be responsible for administering any one group address. To implement such amendments or changes the sponsor of the request would need to contact either the GAA direct, or an agreed representative of the GAA. Only one GAA can be assigned for each group address. A typical centralized database scheme is shown in figure 11.

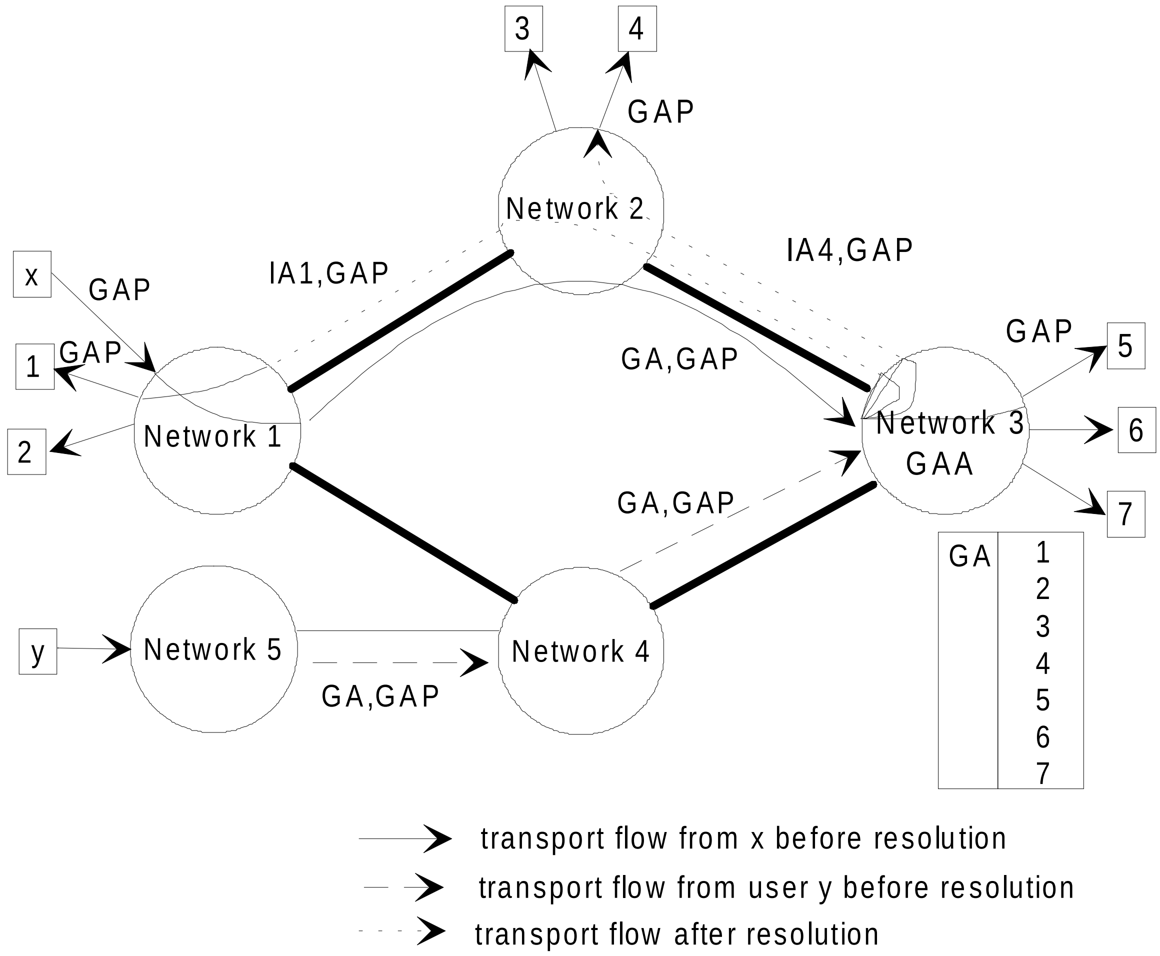

Figure 11: Centralized database approach - group address resolution functions

In this typical example all the information relating to the group address for interconnected Networks 1-3 is held in a central database at Network 3, (the appointed GAA). As shown in the previous example users 1-7 are existing members of a group, accessed by the one group address. User x connected to Network 1 originates a CLNAP_PDU with a GA (GAP) as a destination address. This PDU is routed to Network 3 which contains the GAA for the group address. Network 3 resolves the GA into its own members (5, 6, 7) and sends to each of the external members a packet carrying the individual destination address of the members (e.g. for member 4 an encapsulated packet IA4,GAP) to Network 2. No packet is sent back to originator x. It is a required property of the resolution function to prevent sending back to a given CLAI interface more than one packet. This is achieved as described in annex B in the case where one CLAI corresponds to more than one member of a group. It is a required property of any network connecting a member to prevent sending back after resolution a GAP over the CLAI through which this GAP has originated. This is achieved by checking that a GAP is not to be delivered to the destination CLAI identified by the source address of the GAP CLNAP_PDU. A CLNAP_PDU addressed to GA from user y connected to Network 5 which has no member of this group is routed normally on the destination address. A PDU with group address as destination address will be sent via Network 4 to Network 3 where resolution will take place. The transport mechanism after the resolution works in the same way as above.

NOTE: This approach implies that for transport purposes encapsulation is used to address the individual members once resolution has taken place.

6.3.3.2 Adding a new member to the group

Adding a new member when using a centralized database approach is a simple task, a typical example is shown in figure 12.

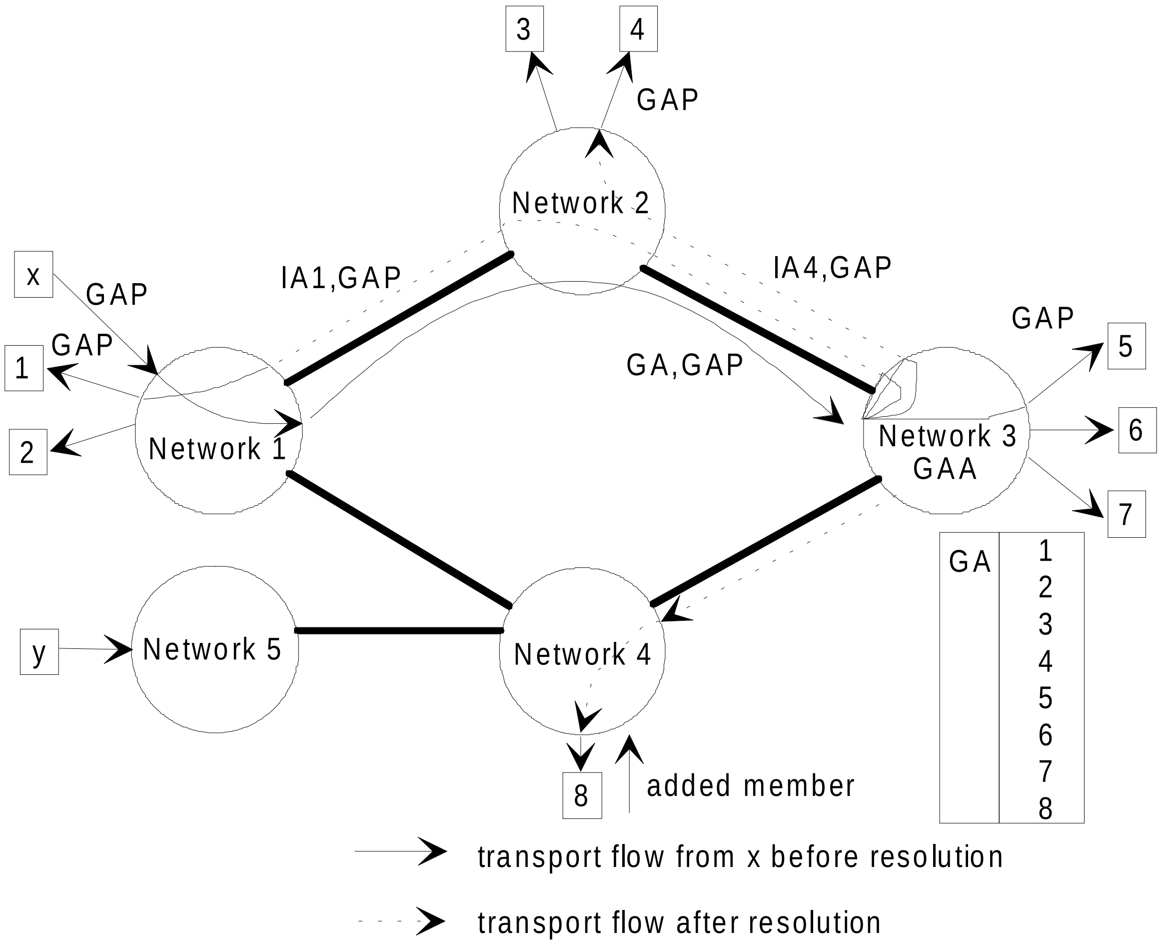

Figure 12: Centralized database approach - adding a new member

For user 8 on Network 4 to be added to the group, the only action necessary is for the relevant information identifying user 8 as a member of the group to be added to GAA database. The encapsulated PDU IA8, GAP between Networks 3 and 4 (for user 8 of this group) is sent in addition to the encapsulated packet IA8 shown in figure 11 when member 8 is added to the group address explained in figure 11.

6.3.4 Centralized approach with Nested Group Address (NGA)

NGA is a concept which could result in significant transport efficiencies through its introduction. For completeness an outline of nested group addressing is given. Administrative aspects related to this approach are detailed in annex B.

Used in conjunction with a centralized database scheme with one group address agent managing each individual group address as described in subclause 4.3.3.1, nested group addressing can be used to avoid sending repetitive information to individual members of a group who are accessed via the same network and introduces the concept of a NGAA. Whilst only one GAA is allowed per group address, one NGAA may be dedicated to serve each separate network, therefore, there may be more than one NGAA per group address. Each network having one or more member of the GA may have the NGAA functionality.

For a given GA it the responsibility of the GAA to determine where the creation of a NGAA, (only one for a given domain), will be advantageous in connecting some members of this GA. This use of only one NGAA in such circumstances would, for a given GAP (GA), prevent more than one NGA, GAP being delivered to a NGAA. In addition it is a required property of the partial resolution function (route via the NGA) when receiving a NGA, GAP not to permit this PDU outside the NGAA domain.

For the above, a NGAA resolution shall:

- prevent sending more that one PDU to a given CLAI of the NGAA;

- prevent sending back the GAP to the CLAI of the domain if this CLAI originated the GAP PDU for which the NGAA is requested to do partial resolution;

- prevent sending the GAP over any CLNI towards another domain.

Figure 13 shows a simple configuration having two NGAAs. Using this example firstly consider the actions required without using nested group addressing.

If users 3 and 4 in Network 2 were included as members of the same group address as users 1, 2, 5, 6 and 7, and information was being sent to them from user 1 in Network 1 then the following actions would be required:

- the group addressed PDU from the originating network (Network 1 in this case) would be passed to the GAA (shown at Network 3);

- group address resolution takes place;

- individually addressed PDUs are copied to users 5, 6 and 7 in Network 3;

- individually addressed PDUs are sent back across the network from the GAA to Network 2 for each individual user identified as a member of that particular group (users 3 and 4 in this case).

Now consider using the nested group addressing capability. The following actions would result:

- the group addressed PDU would be passed to the GAA (shown at Network 3);

- group address resolution takes place;

- individually addressed PDUs are copied to users 5, 6 and 7 in Network 3;

- only one PDU, identified as a "nested" group address is passed between the GAA and Network 2 which is a nominated nested group address server, this is for all users recognized as part of that Nested Group Address 1 (NGA1);

- resolution of the nested group address takes place at Network 2, resulting in the PDU being copied to each individual user who forms part of the group address, in the case shown users 3 and 4. The same considerations apply for the NGA2 (Network 1).

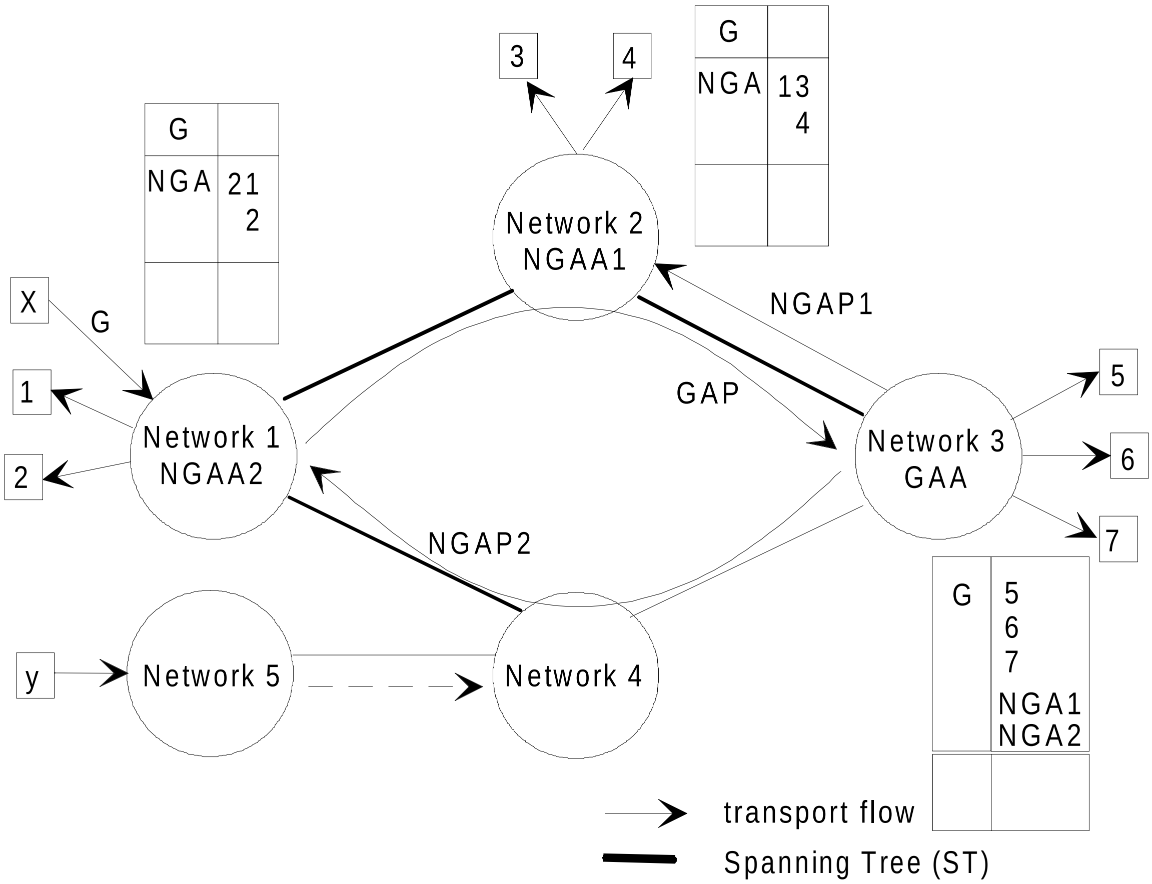

Figure 13: Centralized database approach with nested group addressing

In figure 13 the routeing is not shown. It is assumed that each network knows how to route towards the network identified by a group address or a nested group address. The GAA is solely responsible for ensuring that each NGA receives a nested group address packet NGAP (packet encapsulated with a nested group address NGA, GAP) and of the order in which these nested group address packets are sent from the GAA towards the other networks. It is a local matter to a network to recognize an address as being a nested group address.

If the group address passed to the GAA for resolution also included users in Network 4, then the group address resolution could result in them being associated with the use of another nested group address (NGA3) and Network 4 being another nominated group address server. However, if only one user in Network 4 was a member of that specified group (in this example user 8) then an individually addressed PDU would be sent to Network 4, if no NGAA is created for this network.

Another example of a group address that utilizes three NGAA is shown in figure 14. NGAA 1 servers users 5, 6 and 7 in Network 3, NGAA 2 serves members 1 and 2 in Network 1 and NGAA3 serves users 8, 9 and 10 in Network 4.

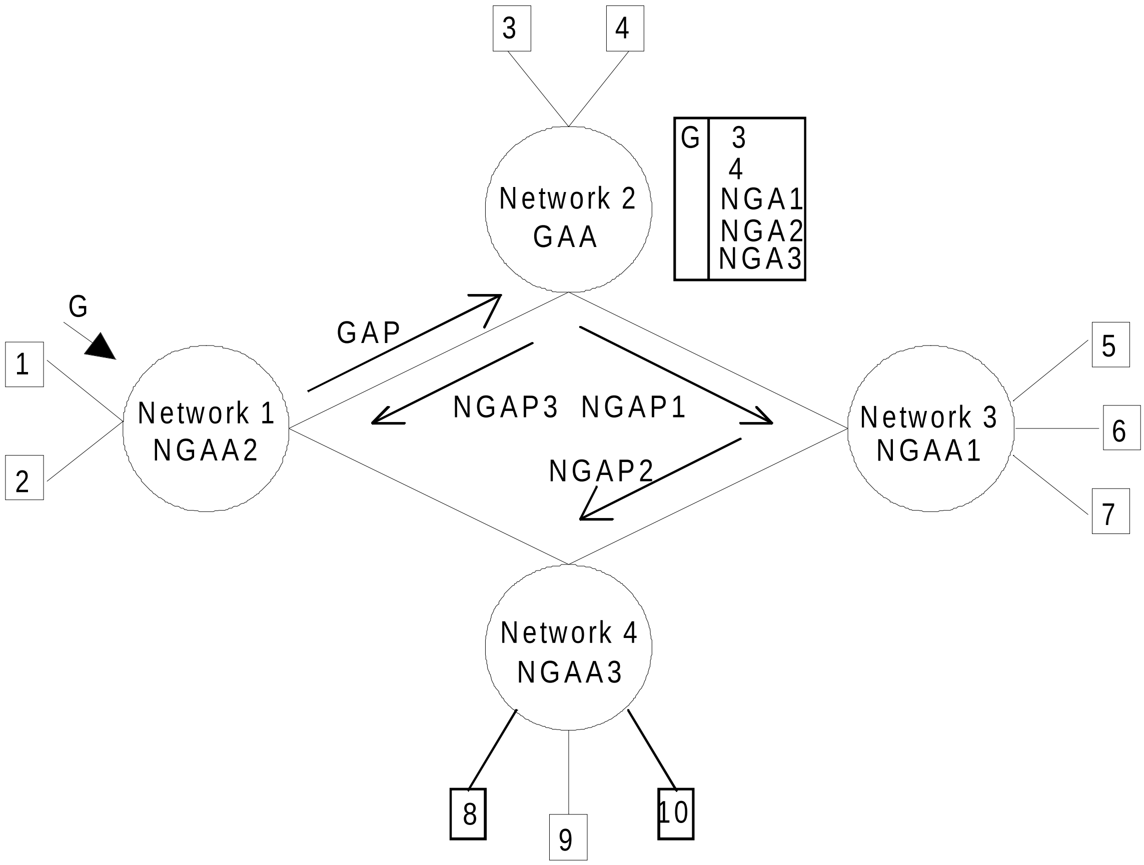

Figure 14: Centralized database approach - use of Nested Group Addresses (NGAs)

In this example it is assumed that working arrangements route all PDUs from Network 2 destined for Network 4 via Network 3. PDUs sent to NGAA3 from the GAA could be delivered for copying to individual addresses for users terminated on Network 4 (users 8, 9 and 10) in the same way.

It would be possible to recognize a nested group address by allocating a new value in the address type field specified in subclause 3.4.3. However, the allocation of such a value would need to be agreed by IEEE (see note), but no requirement to create such a type was found in the above method. A NGAA should be able to recognize a PDU that it has to act upon (copy to a range of individual addresses that it serves) from elements found in the PDU (e.g. addresses) and/or administrative data.

NOTE: An alternative approach of using simple look-up tables for association of NGA's may prove simple to implement, but is for further study.

A number of issues relating to the use of NGA working have to be clarified including:

- routeing;
- methods of identifying NGAs, e.g. sharing group address space;
- additional administration procedures (see annex A);
- additional functionality required when networks receive NGAs (see above);
- responsibilities of intermediate networks.

The specific characteristics related to each approach are listed in annex A.

Annex A: Types of databases

A.1 Distributed databases

- In most cases it results in low network load because only one copy of the data needs to be sent to each of the networks involved.

- Good load distribution, because the multicasting of packets is distributed for each packet.

- Low transit delay because the packets need not be routed to a central multicasting network.

- In some cases provides robustness, e.g. if one network fails the remaining networks (or part of) may still resolve the group address message.

- Relatively easy and rapid to add or delete a new member locally at a network in terms of database modification.

- Difficult to administer and co-ordinate, particularly to groups spread over a number of networks.

- Difficult to have an overall view of the group.

- Spanning Tree (ST) or equivalent will become subject to expansion/reconfiguration as groups expand or change.

- Potentially limits alternative routeing being achieved (see note).

- The overall efficiency of the transport is dependent upon the Spanning Tree (ST) arrangements.

- Tariffing arrangements for customers could also be impacted by the use of Spanning Trees (STs).

NOTE: Routeing aspects are the subject of a separate studywithin STC NA2.

A.2 Replicated databases

- Minimum transit delay because each copy is directly routed towards its destination after resolution within the originating network.

- Robustness, e.g. if one network fails the remaining networks (or part of) may still resolve group address packets.

- Less traffic generated by database queries.

- As the GA related data is copied and then distributed globally local systems can query their own copies.

- Increased network load since the total multicasting occurs at the originating network.

- Delays are incurred in administering changes, because any change needs to be performed in every database, such difficulties increase as the number of networks involved increases, with difficulties in maintaining consistency of up to date information.

A.3 Centralized database

- Overall view of an address group at one network.
- Easier administer and maintain up to date information.
- Facilitates rapid changes.
- The need for consistency of databases is obviated.
- Has the potential to achieve migration and evolution to the nested group addressing methodology.
- Increased transit delay because the packets need to be routed to the central multicasting network.
- Load distribution is limited for a given GA because whatever the source of the group address packets the multicasting of packets is performed by a single network.
- Increased network load because per destination one copy of the packet need to be sent to each individual member of the group in all networks, following GA resolution.
- No robustness in some cases e.g. a single point of failure could result in failure to deliver or administer group address packets.

A.4 Centralized database with Nested Group Address (NGA)

- Provides an efficient mechanism for transport of group address packets.
- Low network load because only one copy of the data needs to be sent to each of the networks involved.
- Good load distribution, because the multicasting of packets is distributed for each packets.
- Relatively easy and rapid to add a new member locally at a network (NGAA) in terms of database modification, when the NGAA is already created.
- Increased transit delay because the packets need to be routed to the central multicasting network.
- No robustness in some cases e.g. a single point of failure could result in failure to deliver or administer group address packets.
- Difficult to have an overall view of the group.
- Tariffing arrangements could be impacted by the use of NGAs.
- Difficult to administer and co-ordinate, particularly to groups spread over a number of networks.

Annex B: OAM flows for GA administration

B.1 OAM when NGA is used

Persons/entities involved:

- CNM;
- CNMA;
- CNM OS;
- GAA;
- NGAA;
- NMS;
- Sponsor.

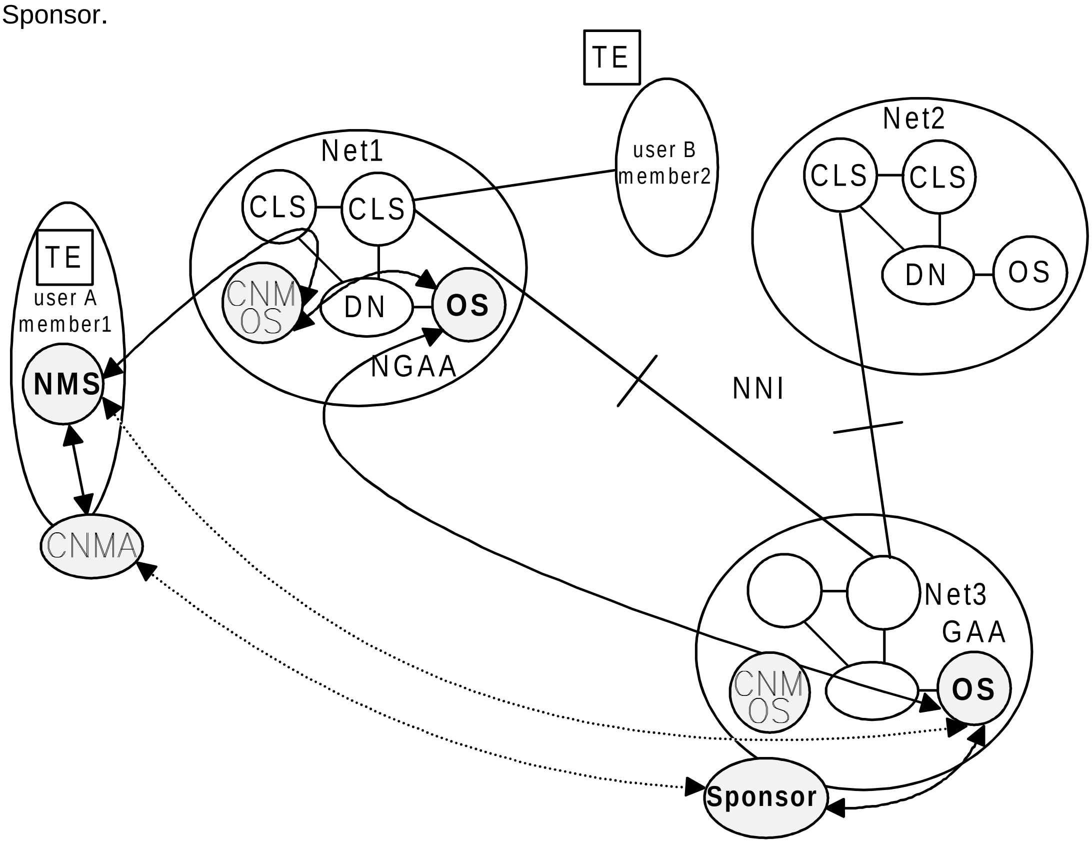

NOTE: Net = Network

Figure B.1: OAM with NGA

The sponsor is a person that deals with the network acting as GAA and with the CNMA who is the customer entity that is responsible for the customer premises associated with the user interface and for the use of the addresses attributed to that interface for a group.

CNM is a possible service, not currently defined for CBDS, that, when provided by the network to a customer, allows this customer to interrogate the network about its subscription parameters, customer traffic, protocol performance, and possibly modify its subscription parameters.

Since sponsor and CNMA are not network entities, the results of their action on the OS/NMS of their respective networks will be described if necessary in terms of actions on the OS/NMS of their respective networks.

The GAA or NGAA is represented by the OS of the associated network.

Nothing prevents in principle a CNMA to dialogue via the NMS with the OS of the GAA (the NGAA is not seen by the user), but this will not be described for simplicity.

The relations can be then drawn in the following way:

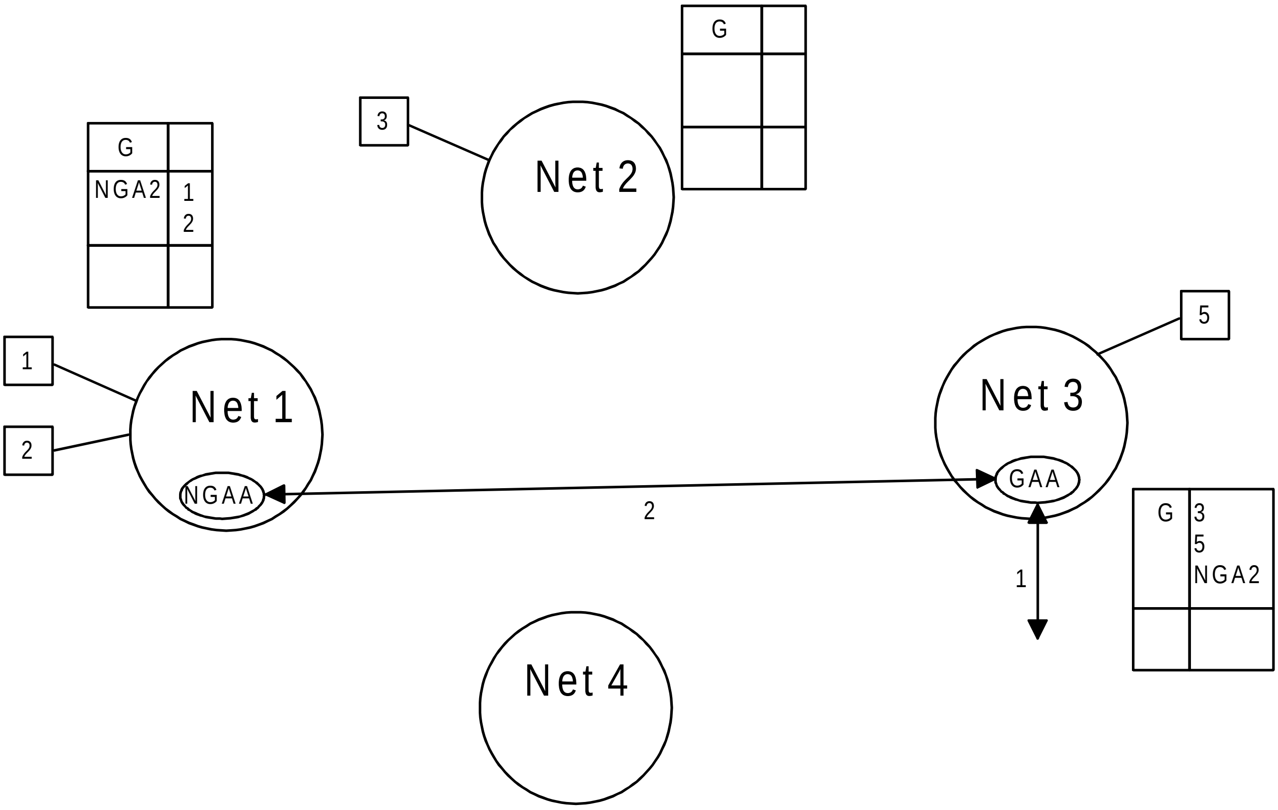

NOTE: Net = Network

Figure B.2: Relation GAA-NGAA example

The way to transport the flows shown between the entities is not specified; its could be achieved via different means e.g. via OS-OS protocols for GAA-NGAA or CNMOS-NGAA relations, SNMP over CBDS for NMS-CNMOS, or simply mail, fax, telephone, etc. The issues of security/confidentiality are not addressed here but are surely relevant to the administration of GAs.

Main actions to be considered are:

- creation of a group;
- creation of a nested group address;
- modification of a group by addition/deletion of a member;
- enquiry on a group;
- deletion of a group;
- enquiry on a member (for further study).

B.1.1 Sponsor

It is the person/entity that is in charge of the group vis-à-vis the outside world, from its creation until its deletion.

Its maintains per GA that it is responsible for, the group address GA, the list of all members, the individual address IA of each member of the group, the list of CNMA associated with each member, the status of the group.

It requests from a network that will be the GAA for this group the creation/deletion/modification as well as enquiry of a group.

It needs from the CNMA the individual address of the member handled by this CNMA and the availability of this member.

It informs the CNMAs of the availability of the group and of all modifications of the composition of the group.

B.1.2 CNMA

It is the entity/person that is primarily responsible of the use of the addresses assigned to the user-network interface for group addressing.

It asks the sponsor:

- to add/delete an IA assigned to the interface as a member of a given group; and
- to provide to it information about the group.

It maintains and controls the list of GAs identifying the interface, and the list of IAs assigned to the interface and that identify each of these GAs. To do this it may need to control the service features limits supported by the connecting network (e.g. Net 1 for user A in figure B.1) possibly via CNM OS.

The user may use group address screens as a service feature (see ETS 300 217-4 [4]); if group address destination screen screens are set with the allowed option, it is necessary to include/delete the GA from these screens when creating/modifying/deleting a group in order to allow communication with this GA as a destination address. The CNMA accomplishes this, possibly via the NMS-CNM OS if the capability customer network management exists, or via service order to the connecting network (e.g. Net 1 in figure B.1) in the contrary.

The CNMA needs also to take appropriate actions in the user premises in order to allow GAPs with this GA as a destination address to be sent/received by the user premises. The CNMA may need to control that only one IA is assigned to the interface for a given group since the GAA may not be aware of the identity of the interface if it is a distant network, i.e. not the connecting network. (This is only needed if no NGA is created), this being required in order not to deliver more than one copy of the GAP to the user interface.

A CNMA could enquire about the group to the GAA but this requires further study since there can be many - up to the size of the group, e.g. 128-CNMAs inquiring to a GAA and it may cause problems of security and complexity to the GAA.

B.2 Interaction modelling

Interactions between persons, e.g. the sponsor and the CNMA are not modelled. Interactions between the persons sponsor/CNMA and the appropriate entities GAA/NMS are modelled as possible man-machine interactions on the GAA OS/NMS but are not intended to be specifications of the actual interactions inside the network since the person may not have direct access to the other entity. Interactions in the user premises are not modelled since it is outside the scope of this ETR.

B.2.1 Sponsor-GAA interaction

B.2.1.1 Creation of a group

Sponsor-->GAA (Create_GA_Req)

action = request, function = create, request reference, list of the IAs of the members

Inclusion of a list of operators concerned by the group is for further study.

GAA-->Sponsor (Create_GA_Ack.)

action = response, request reference, status, GA, time when GA will be operational, where status indicates if the request has been successful or not with possible indication of a cause in case of failure.

B.2.1.2 Deletion of a group

Sponsor-->GAA (Delete_GA_Req)

action = request function = delete, request reference, GA

GAA-->Sponsor (Delete_GA_Ack.)

action = response, request reference, status, GA, (time when GA is no longer operational)

Deletion of more than one group handled by the same network could be envisaged by extending the list of GAs and the status provided in the request/response; multiple responses, one per GA could be a possibility.

B.2.1.3 Group enquiry

sponsor-->GAA (Enquiry_GA_Req)

action = request, function = enquiry, request reference, GA (or list of GAs)

GAA-->sponsor (Enquity_GA_Ack)

action = response, request reference, status, GA, status of group, list of IAs of members

B.2.1.4 Group modification

sponsor-->GAA (Modify_GA_Req)

action = request, function = add/delete, request reference, GA, list of their IA of members

Alternatively, one could provide a list of IA of members to add, a list of operators to add, a list of IAs to delete and a list of operators to delete.

GAA-->sponsor (Modify_GA_Ack)

action = response, reference, GA, Status, time when action done

B.2.2 GAA-NGAA interaction

B.2.2.1 Group creation

GAA-->NGAA

action = request, function = create, request reference, GA, list of IAs for NGAA members

NGAA-->GAA

action = response, request reference, status, GA, NGA, time when operational

The NGA identifies the network (NDC of E164) that acts as the NGAA for the GA, and the set of IAs of the members of this GA served directly by this NGAA network.

B.2.2.2 Deletion of a group

GAA-->NGAA

action = request, function = deletion, request reference, GA, NGA

NGAA-->GAA

action = response, reference, status, GA, NGA, (time when NGA will no longer be operational).

Deletion of more than one nested group is for further study, but in any case there is always a single NGA per NGAA associated to a GA.

B.2.2.3 Group enquiry

GAA-->NGAA

action = request, function = enquiry, request reference, GA, NGA

NGAA-->GAA

action = response, request reference, status, GA, NGA, status of group, list of IAs, (status of IAs)

This enquiry may be done periodically by the GAA or upon a request from the sponsor to the GAA (may be acting upon a complaint from a CNMA that packets are not delivered to the user).

B.2.2.4 Group modification

GAA-->NGAA

action = request, function = modification, type = add/delete, request reference, GA, NGA, list of IAs

Alternatively, a list of IAs to add and a list of IAs to delete combined with a list of operators to add and delete could be introduced.

NGAA-->GAA

action = response, reference, GA, NGA, status, time when action done

B.2.3 CNM OS-NMS interactions

They are relevant only if the customer network management capability is supported by the user and the connecting network and that in addition the user can through this capability modify its service features. Otherwise, this should be replaced by relations between the CNMA and the service order entity of the connecting network (which will provide the man machine communication to the OS).

The main capabilities involved in customer network management are retrieve/modify subscriber profile information (subscriber address and address screening tables, network group address table related to that interface.), and receive event notification (subscriber entry change, group address entry change, group screen change).

These capabilities assume that interactions exist between the connecting node and the CNM OS, and between the OS and the CNM OS in the case the CNM OS and the OS are separate entities. Interactions between the OS and the connecting node are needed regardless customer network management being provided or the respective location of the CNM OS and the OS.

All these interactions will not be modelled for the sake of simplicity.

B.3 Group creation example

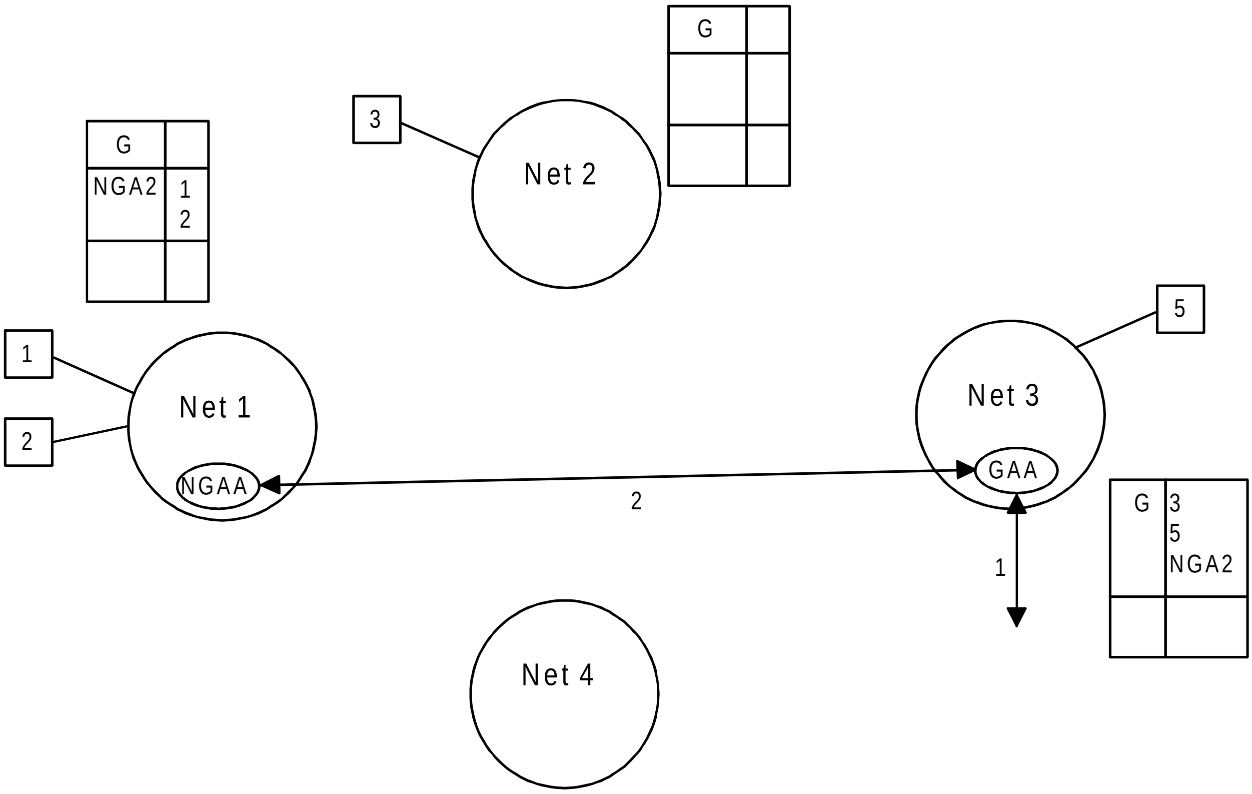

NOTE: Net = Network

Figure B.3: Group creation example

Once the sponsor has the list of IA of the intended members of the group to create, a request 1 is issued to the GAA in Network 3.

The GAA assigns a GA to the new group, adds members 3 and 5 to the group and decides to request a NGA from Network 1 which serves members 1 and 2. No NGA is requested from Network 2 that only connects member 3.

Network 1 receives a request 2 from the GAA and assigns a NGA2 encompassing members 1 and 2 and responds to the GAA with a confirmation of NGA2 and the time "T" at which the NGA will be provisioned.

The sponsor will confirm to each CNMA (users 1, 2, 3 and 5) the creation of GA at time "T" and the IA member of the group which is under its responsibility. Each CNMA will do what is necessary so that GA becomes operational in its premises and that the connecting network screens are updated if needed.

History

Document history	
May 1997	First Edition

ISBN 2-7437-0550-7
Dépôt légal : Mai 1997