INTERNATIONAL TELECOMMUNICATION UNION

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

X.601

(03/2000)

SERIES X: DATA NETWORKS AND OPEN SYSTEM COMMUNICATIONS

OSI networking and system aspects – Networking

Multi-peer communications framework

ITU-T Recommendation X.601

(Previously CCITT Recommendation)

ITU-T X-SERIES RECOMMENDATIONS

DATA NETWORKS AND OPEN SYSTEM COMMUNICATIONS

PUBLIC DATA NETWORKS	
Services and facilities	X.1–X.19
Interfaces	X.20–X.49
Transmission, signalling and switching	X.50–X.89
Network aspects	X.90–X.149
Maintenance	X.150–X.179
Administrative arrangements	X.180–X.199
OPEN SYSTEMS INTERCONNECTION	
Model and notation	X.200–X.209
Service definitions	X.210–X.219
Connection-mode protocol specifications	X.220–X.229
Connectionless-mode protocol specifications	X.230–X.239
PICS proformas	X.240–X.259
Protocol Identification	X.260–X.269
Security Protocols	X.270–X.279
Layer Managed Objects	X.280–X.289
Conformance testing	X.290–X.299
INTERWORKING BETWEEN NETWORKS	
General	X.300–X.349
Satellite data transmission systems	X.350–X.399
MESSAGE HANDLING SYSTEMS	X.400–X.499
DIRECTORY	X.500–X.599
OSI NETWORKING AND SYSTEM ASPECTS	
Networking	**X.600–X.629**
Efficiency	X.630–X.639
Quality of service	X.640–X.649
Naming, Addressing and Registration	X.650–X.679
Abstract Syntax Notation One (ASN.1)	X.680–X.699
OSI MANAGEMENT	
Systems Management framework and architecture	X.700–X.709
Management Communication Service and Protocol	X.710–X.719
Structure of Management Information	X.720–X.729
Management functions and ODMA functions	X.730–X.799
SECURITY	X.800–X.849
OSI APPLICATIONS	
Commitment, Concurrency and Recovery	X.850–X.859
Transaction processing	X.860–X.879
Remote operations	X.880–X.899
OPEN DISTRIBUTED PROCESSING	X.900–X.999

For further details, please refer to ITU-T List of Recommendations.

ITU-T RECOMMENDATIONS SERIES

Series A	Organization of the work of the ITU-T
Series B	Means of expression: definitions, symbols, classification
Series C	General telecommunication statistics
Series D	General tariff principles
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Construction, installation and protection of cables and other elements of outside plant
Series M	TMN and network maintenance: international transmission systems, telephone circuits, telegraphy, facsimile and leased circuits
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	**Data networks and open system communications**
Series Y	Global information infrastructure and Internet protocol aspects
Series Z	Languages and general software aspects for telecommunication systems

18573

Printed in Switzerland

Geneva, 2000

ITU-T RECOMMENDATION X.601

MULTI-PEER COMMUNICATIONS FRAMEWORK

Summary

This Recommendation provides the basic framework to specify services and protocols for multi-peer communications. It also defines the basic concepts of group and various aspects of group communication, which are needed to specify specific services and protocols for multi-peer communications.

Source

ITU-T Recommendation X.601 was prepared by ITU-T Study Group 7 (1997-2000) and was approved under the WTSC Resolution No. 1 procedure on 31 March 2000.

FOREWORD

ITU (International Telecommunication Union) is the United Nations Specialized Agency in the field of telecommunications. The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of the ITU. The ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Conference (WTSC), which meets every four years, establishes the topics for study by the ITU-T Study Groups which, in their turn, produce Recommendations on these topics.

The approval of Recommendations by the Members of the ITU-T is covered by the procedure laid down in WTSC Resolution No. 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

INTELLECTUAL PROPERTY RIGHTS

The ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. The ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, the ITU had not received notice of intellectual property, protected by patents, which may be required to implement this Recommendation. However, implementors are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database.

CONTENTS

MULTI-PEER COMMUNICATIONS FRAMEWORK

1 Scope

This Recommendation addresses the basic concepts needed to specify services and protocols for multi-peer communications. It defines the related terminology, and proposes a framework for the future development of multi-peer services and protocols.

Multi-peer is limited to the service view of communication among more than two participants, and multi-peer communication presents a mode of operation that supports exchanges between more than two service-users.

Multicast data transmission is defined as a transmission of the same data unit from one sender to a set of receivers in a single invocation of a service.

2 Normative references

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; all users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of currently valid ITU-T Recommendations is regularly published.

2.1 Identical Recommendations | International Standards

– ITU-T Recommendation X.200 (1994) | ISO/IEC 7498-1:1994, *Information technology – Open Systems Interconnection – Basic Reference Model: the Basic Model.*

– ITU-T Recommendation X.214 (1995) | ISO/IEC 8072:1996, *Information technology – Open Systems Interconnection – Transport service definition.*

– ITU-T Recommendation X.605 (1998) | ISO/IEC 13252:1999, *Information technology – Enhanced Communications Transport Service definition.*

– ITU-T Recommendation X.641 (1997) | ISO/IEC 13236:1998, *Information technology – Quality of Service – Framework.*

3 Definitions and abbreviations

This Recommendation defines the following terms:

3.1 Terms defined in the OSI Basic Reference Model

This Recommendation builds on the concepts developed in ITU-T Rec. X.200 | ISO/IEC 7498-1 and makes use of the following terms defined in that Recommendation:

a) (N)-entity

b) (N)-layer

c) (N)-protocol

d) (N)-service

e) (N)-service-access-point

f) (N)-service user

g) (N)-subsystem

3.2 Abbreviations defined in the OSI Basic Reference Model

This Recommendation builds on the concepts developed in ITU-T Rec. X.200 | ISO/IEC 7498-1 and makes use of the following abbreviations defined in that Recommendation:

a) (N)-SAP

b) SDU

c) PDU

d) CEP

3.3 Abbreviations defined in this Recommendation

The terms and definitions are found in the subclauses in which they are defined.

4 Notation

Layers are described in ITU-T Rec. X.200 | ISO/IEC 7498-1. And (N)-, (N+1)-, and (N–1)- notation is used to identify and relate adjacent layers:

– (N)-layer: any specified layer

– (N+1)-layer: the next higher layer

– (N–1)-layer: the next lower layer

This notation is also used for other concepts in this Recommendation which are related to these layers, for example (N)-protocol, (N+1)-service.

When referring to a layer by name, the (N)-, (N+1)-, (N–1)- are replaced by the names of the layer, for example transport-protocol, network-service, etc.

5 Description of multi-peer

5.1 Definitions

5.1.1 multi-peer: A service view of communication among more than two participants.

5.1.2 multi-peer communication: A mode of operation that supports exchanges between more than two service-users.

5.1.3 multicast transmission: A transmission of the same data unit from a single source to multiple destinations in a single invocation of a service.

5.1.4 (N)-group-connection (GC): A multicast connection established among (N)-service users for the purpose of transferring data, which is established by the (N)-layer.

5.1.5 (N)-group-connection-end-point (GCEP): A terminator at one end of an (N)-group-connection.

5.1.6 (N)-group-association (GA): A cooperative relationship among (N)-entity-invocations. In an (N)-connection-mode service, the establishment of an (N)-group-association shall be a set of (N)-group-connections, which has a cooperative relationship among the entities in the next higher or above layers. An (N)-GA may contain one or more (N)-GCs in it.

5.1.7 (N)-group-association-end-point (GAEP): A terminator at one end of an (N)-group-association within an end system.

5.2 Description

For information to be exchanged among (N)-service users, an (N)-group-connection is established among them in the (N)-layer using an (N)-protocol. A set of (N)-group-connections may be formed to an (N)-group-association by the entities in the next higher or above layers.

The rules and formats of an (N)-protocol are instantiated in an (N)-subsystem by an (N)-entity. An (N)-entity may support one or more (N)-protocols.

(N)-entities when supporting multi-peer communication maintain the binding of (N)-group-connections to the appropriate (N)-service users at the appropriate (N)-SAPs.

6 Group descriptions

The concepts of groups presented here are applied to modeling components of multi-peer communication, which are: group application, group session, group association, and group connection.

6.1 Definitions

The concept of group allows defining a set of entities as a virtual single entity. One of the main reasons to create a group, is the possibility to name and address all group members with a single name, a single address, or unified characteristics of group communications, called respectively, group-name, group-address, or group-characteristics.

6.1.1 (N)-group: A set of (N)-service users using (N)-service.

6.1.2 (N)-multicast group: A set of (N)-service users that abide by appropriate (N)-group-membership criteria, or a set of rules for belonging to a group able to utilize (N)-multicast services. To each (N)-multicast group, an (N)-group name is assigned. The (N)-group-name and the rule defining the (N)-multicast group are known, but it may not be feasible to determine all (N)-service users that satisfy the rule.

Note that (N)-multicast service can be provided by (N)-entities or (N–1) services.

6.1.3 (N)-registered group: A set of (N)-multicast group members, which has expressed implicitly or explicitly to the (N)-group manager that they have the intention to be a member of (N)-enrolled-group. (N)-Group manager collects peer names and addresses of (N)-registered group members. Though it may not be possible for any (N)-service user to determine the members of the (N)-registered group, it may be possible for a certain (N)-service user to determine the members of the (N)-registered group and keep the list of those members.

6.1.4 (N)-enrolled group: A set of (N)-registered group members to which an (N)-group address is assigned. An (N)-enrolled group is actually in the position to be reached by means of an (N)-group address. A member of an (N)-enrolled group can also be reached by its individual address. In connection mode service, an (N)-group-connection is established among (N)-enrolled group members.

6.1.5 (N)-active group: A set of (N)-enrolled group members that has entered the multicast data transfer phase, satisfying the group-characteristics requirements.

6.2 Descriptions

The use of the generic term, group and the qualified terms such as multicast-group defines a hierarchy of group types.

The most general group type is the (N)-multicast group and consists of the (N)-service users that abide by the (N)-group-membership criteria. The criteria may be defined by means of a set of rules or by more arbitrary means such as a list. The (N)-group-membership criteria may include, (N)-group-name(s), the list(s) of (N)-group members addresses, proposed (N)-group-characteristics, and so on, depending on the hierarchy in which the group is located. An (N)-multicast group is identified by an (N)-group name.

The (N)-registered group is a proper subset of the (N)-multicast group, and consists of (N)-service users that have announced their ability to participate in (N)-group communications. After this registration operation, a member of an (N)-registered group should be available to the group members to be used for multi-peer communication. This is performed by the enrollment phase, in which an (N)-group address is assigned to the (N)-registered group, and it creates an (N)-enrolled group. Therefore, an (N)-enrolled group can be defined as a set of members that belongs to an (N)-multicast group and has successfully completed the registration and enrollment. The member of an (N)-enrolled group is actually in the position to be reached by the means of one (N)-group-address. A member of an (N)-enrolled group can also be reached by its individual address. At this point, the (N)-enrolled group:

– consists of the set of (N)-service users that may participate in an (N)-connection;

– is identified by the (N)-group name(s) and (N)-group address(es);

– is a proper subset of the (N)-multicast group and an (N)-registered group.

When an (N)-service user wishes to participate in an instance of communication by sending or receiving data over an (N)-group-connection, it joins the (N)-active group corresponding to the (N)-group connection. An (N)-active group is a subset of the (N)-enrolled group and consists of those (N)-service users.

The hierarchy of group types is depicted in Figure 1.

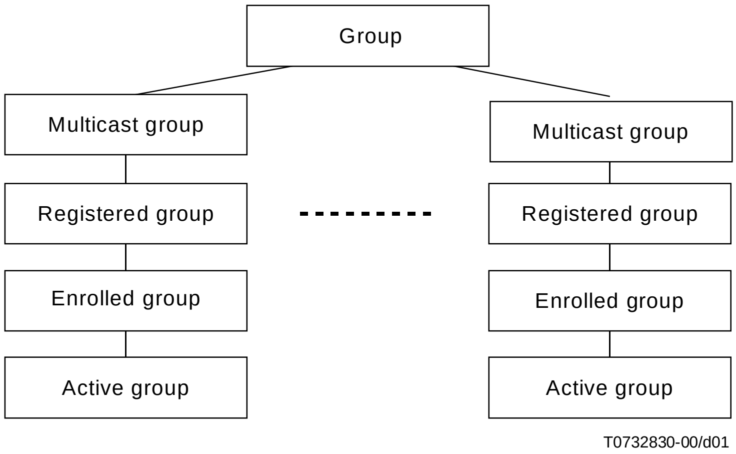

Figure 1/X.601 – Hierarchy of group types

6.3 Group characteristics

6.3.1 Population characteristics

Three distinctions are described here: static vs. dynamic population, known vs. unknown population, and open vs. closed population.

Static vs. dynamic population: a *static group* is one in which the population of the group does not change. A *dynamic group* is one in which the population of the group may change.

Known vs. unknown population: a *known group* is a group for which there is a way to unambiguously determine the complete set of members. Not all members may be aware of the group population. Within a known group, there may be one of two sub-categories of population knowledge;

– complete-known population where every member may know the individual name or address of every member of the enrolled-group;

– partial-known population where a subset of members may know the individual name or address of every member of the enrolled-group.

Unknown group is a group for which it is not possible to unambiguously determine the group membership. Within an unknown group, some members may be known by all or part of the whole population. In that case, these members are called known members.

Open vs. closed population: an *open group* is one where a non-registered member of the group may take part in active group. A *closed group* is one where only registered members may take part in active group.

6.3.2 Communication discipline

Communication discipline describes the allowed group behavior of the sender(s) and receivers on a multicast group with respect to the transfer of data.

Send only vs. receive only: within a group, a sender(s) may originate multicast transmissions that are not associated with any response(s) from recipients. Within the same group, receivers may only receive transmissions.

Send/Receive: Within a group all members may either send or receive messages.

6.3.3 Dialog control

Centralized vs. decentralized: a *centralized group* is one in which a single (designated) member is permitted to send, and all other members are permitted to receive. Depending on the communications discipline, it may or may not be possible for receivers in a centralized group to respond to communications originated by the sender.

The designation of a member as sender may be dynamic and change from one member to another as long as only a single member is allowed to send at any given time.

A *decentralized group* is one in which any member is permitted to either send or receive (or both).

NOTE – It is possible for restrictions to be placed on individual members of the active group on whether they are permitted to send data. That is, in an *decentralized group*, only a subset of the active members may be authorized to send data.

6.3.4 Concurrency control

Controlled vs. uncontrolled: A *controlled group* is one in which that only senders with permits can transmit data, while an *uncontrolled group* is one in which all senders may transmit data concurrently.

6.3.5 Reliability characteristics

A new feature of reliable data transfer that appears on a multicast data transmission is the degree of reliability. Reliability spans a range from *fully reliable* (in-order error-free delivery is guaranteed to all members of an active group) to *not reliable* (no guarantees concerning the order of delivery or presence of errors are made). Reliability may also be expressed with respect to the AGI characteristics, which allows for variations such as "fully reliable with respect to a specified subset or an unspecified quorum of the members of the group, don't care beyond that", etc. This can be defined as a *k-reliable* wherein k ranging from zero to the total active group size, but fixed. In the case of *k-reliable*, data-units sent from a sender should be received by at least k active receivers.

A special kind of fully reliable data transfer is the *synchronized transfer* where a data unit is delivered either to all active receivers or to none of them at each moment. An *unsynchronized transfer* does not have that feature.

6.4 Naming and addressing

6.4.1 Definitions

6.4.1.1 (N)-group-name: A name that is used to identify a set of (N)-service users, or (N+1)-entities that identifies members of an (N)-group.

6.4.1.2 (N)-group-address: An address that is used to identify a set of (N)-addresses that identifies members of an (N)-group.

6.4.2 Naming

An (N)-multicast group can be identified by its (N)-group-name. The (N)-group-name must be unambiguous within an open system environment to ensure that entities registering for a particular (N)-multicast group can unambiguously identify which multicast group it wishes to join. The (N)-group-name identifies only the multicast group, and identifies neither (N)-service users, (N+1)-entities nor the mapping of (N)-service users to the group and addresses. The (N)-group-name is used for the purposes of identifying the existence of the (N)-multicast group to the (N)-service users in the open system environment. An (N)-group-name is assigned and managed by an appropriate naming authority under the auspices of system management.

6.4.3 Addressing

An (N)-group-address is an unambiguous name within an open system environment that is used to identify the set of (N)-addresses that identify members of a group. When an (N)-group-address is used to send a data unit, all members of the list receive the data unit. An (N)-group-address may stand for a list of (N)-addresses of (N)-group members.

In order to identify each (N)-group association, an (N)-GA-identifier is associated with each (N)-GA. It serves to identify the related state information associated with each participant of a group association. In order to identify the group connections within a given GA, an (N)-GC-identifier is associated with each (N)-group-connection of a GA. It serves to identify the related state information associated with each participant of a group association.

Each terminal of an (N)-GA and (N)-GC is called an (N)-GA-endpoint (GAEP) and (N)-GC-endpoint (GCEP), respectively. Each (N)-GAEP is identified by an (N)-GAEP-identifier, and each (N)-GCEP is identified by an (N)-GCEP-identifier.

6.5 Active Group Integrity (AGI)

The *active group integrity (AGI)* specifies conditions on the active group membership. In other words, AGI describes the conditions that must be present in order for the data transfer to take place. The AGI relates to the characteristics of the active group. AGI is not an attribute of the individual active group members, but an attribute of the group. If the AGI is no longer met during the data transfer phase, the data transfer may be either terminated (*hard AGI*) or suspended (*soft AGI*) until conditions are again satisfied according to the *AGI policy*. A connectionless mode of communication has no AGI condition, because a receiver is identified as active when it participates in the data transfer phase. The *AGI population characteristics* can be one or more of the following conditions:

– Minimum: condition that specifies the minimum number of enrolled group members required to be presented in the active group;

– Quorum: condition that specifies the majority of enrolled group members required to be presented in the active group;

– Maximum: condition that specifies the maximum number of enrolled group members that can be allowed in the active group;

– Mandatory: condition that specifies the selected enrolled group members required to be presented in the active group;

– Atomic: condition that specifies all of the enrolled group members required to be presented in the active group.

6.6 Ordering

Ordering is concerned with the following two aspects:

– how PDUs of a single sender are presented to the receivers;

– how a single receiver gets PDUs the sender(s).

In the case of a single sender, ordering, if needed, ensures that the data units generated by the sender are delivered to each receiver in the active group in the same order as they were sent. In the case of multiple senders, ordering determines the relative sequencing of data received from multiple senders. The ordering relationship defines the arrangement or interleaving of data from the multiple senders.

The ordering relationship can be: no, local, partial, causal, or total. Note that when there are only two participants in the active group, local ordering, causal ordering, and total ordering are the same.

The properties of ordering apply at the service level and at the protocol level. At the service level, the service provider may be required to provide guarantees regarding the order in which SDUs are delivered to the receiving TS-users. At the protocol level, PDUs are ordered, or reordered to achieve the ordering property required by the service.

The following notation is used to describe the ordering relationships:

S_i(m): Local event of sending data-unit m at site i (=1,2,...,N);

A_j(m): Local event of accepting the data-unit m at a site j;

A ---> B: An event A happened before an event B;

A ==> B: If A is TRUE, then B is to be TRUE;

A =/=> B: If A is TRUE, then B does not need to be TRUE.

6.6.1 No ordering

The service provider does not guarantee any relationship between data-units sent from a single sender or from multiple senders.

Notation : S_p(m) ---> S_q(m') =/=> A_i(m) ---> A_i(m')

for all p,q,i and for all (m,m') pairs.

6.6.2 Local ordering

The data-units, generated by a particular sender, are delivered to all the receivers in the same order in which they were sent. Local ordering does not establish any ordering relationship among data-units sent by different senders.

Notation : S_p(m) ---> S_p(m') ==> A_i(m) ---> A_i(m')

for all p,i and for all (m,m') pairs.

But, the following constraint also applies:

Notation: S_p(m) --->S_p(m') ==> A_i(m) ---> A_i(m')

for any given p,i pair and for all (m,m') pairs.

6.6.3 Partial ordering

The data-units, generated by all the senders, are delivered to each receiver according to an arbitrary ordering rule. If the data-units are ordered according to a rule applicable to all the receivers, then each receiver receives the data-units generated by all the senders in the same order. If the data-units are ordered according to a rule determined by each receiver, then each receiver may receive the data-units in a different order.

Notation : If the arbitrary ordering rule is set by the service provider for all receivers,

Then S_p(m) ---> S_q(m') =/=> A_i(m) ---> A_i(m')

for all i but for some p,q and for all (m,m') pairs.

or

If the arbitrary ordering rule is set independently by each receiver,

Then S_p(m) ---> S_q(m') =/=> A_i(m) ---> A_i(m')

for some p,q,i and for some (m,m') pairs.

6.6.4 Causal ordering

Causal ordering orders the data-units generated by all the senders according to the causal dependence relationship among the sending events. A causal dependence relationship is established between two sending events, A and B, if the following applies:

– A happens before B if A and B are sending events generated by the same sender and A is sent before B;

– A happens before B if A and B are sending events generated by two different senders and the data-unit generated by the event A by one sender is received by the other sender before it generates the event B.

A causal dependence relationship is established among more than two sending events if it is established that A happens before B and that B happens before C, then it can be established that A happens before C. A causal dependence relationship cannot be established between the two sending events A and C if there is no possibility to establish that A happens before B and that B happens before C.

Notation : (S_p(m) ---> A_q(m) ---> S_q(m')) or (S_q(m) --> S_q(m')) ==> A_i(m) ---> A_i(m')

for all p,q,i and for all (m,m') pairs

6.6.5 Total ordering

The data-units, generated by all the senders, are delivered to each receiver in the same order. Every receiver sees all data-units from all senders in exactly the same order.

Notation : S_p(m) ---> S_q(m') ==> A_i(m) ---> A_i(m')

for all p,q,i and all (m,m') pairs

6.7 Synchronization

Synchronization on data-units transmitted from several sources may be required in multimedia systems. This requirement leads to the definition of a new facility called orchestration that allows the sychronization of data-units transmitted within different connections.

7 Model of multi-peer communications

7.1 Group Association (GA)

A cooperative relationship among two or more (N)-enrolled groups is needed so that each member of an (N)-enrolled group takes part in a group communication with other members of its own (N)-enrolled group while it keeps a cooperative relation with other (N)-enrolled group members.

An (N)-group-association, or (N)-GA is an association established between (N)-enrolled-groups, and possibly (N)-group members, not belonging to any (N)-enrolled group if it is not restricted by the (N)-group membership characteristics, for the purpose of transferring data.

The members in an (N)-GA are said to be *participants* of the (N)-GA. The concept of participant is not restricted to (N)-enrolled group members. This means that any (N)-service users that is taking part in an (N)-GA is considered as a participant of that (N)-GA whether it is an enrolled member or not, if it is not explicitly restricted by the use of group membership characteristics.

The members who take part in an (N)-GA are called (N)-GA-participant.

In each (N)-GA, there may be multiple concurrent (N)-enrolled groups. Each (N)-GA at an end system is identified by the (N)-GAEP.

The concept of group association covers both connection-oriented (CO) and connectionless (CL) transmission modes. The establishment of an (N)-GA may lead to the use of an (N)-CO-mode service and subsequently the creation of (N)-GC(s), and/or to the use of an (N)-CL-mode service.

7.1.1 GA topology integrity

There is no a priori restriction on the topology of a group association. Hence, a group association may look like Figure 2, where each set of identical oriented lines represents a data flow. From this, it is clear that, for a given group-association, many different topologies are possible.

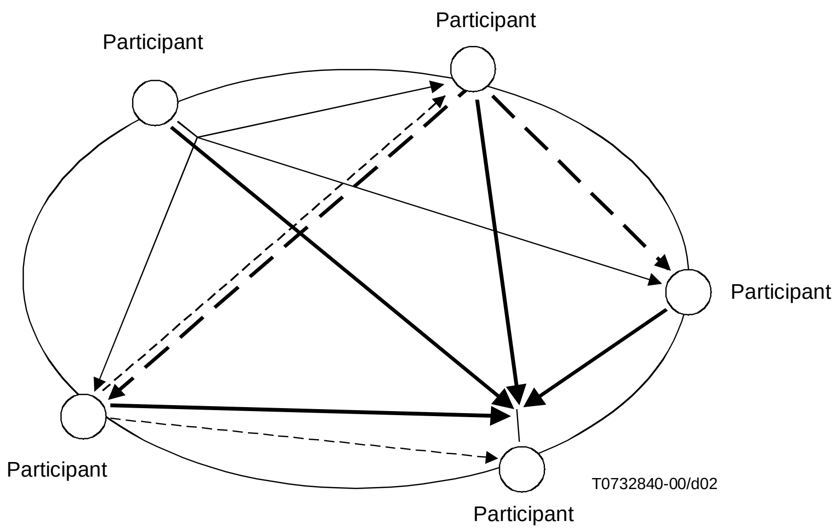

Figure 2/X.601 – Example of GA topology

A complex group association topology, as shown in Figure 2, can always be seen as a set of group-connections (GCs). In the simplest case, there is only one possible group-connection defined within a group-association (GA). As an example, in the case of a multimedia communication, a GA may be composed of a set of group-connections, each of them conveying different traffic (e.g. voice, data, and video). The relationship between group-association and group-connection is illustrated in Figure 3.

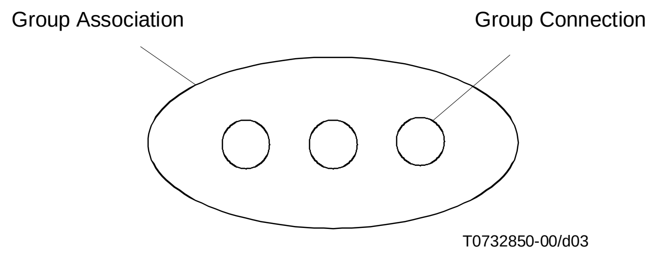

Figure 3/X.601 – Relationship between GA and GCs

As with the active group, conditions may be specified on the general topology of a group association. These conditions, which may be the same as AGI population characteristics such as minimum, quorum, maximum, mandatory and atomic, are called the GA topology integrity (ATI).

The ATI is composed of conditions on the whole group association. If the conditions on the whole group association are no longer satisfied, the group association is either released or suspended until the conditions are again satisfied.

As an example of the use of the ATI concept, ATI can be used in a multimedia communication to impose the release of the group association when the group-connection related to the required traffic (voice for instance) fails. In that case, the ATI is: a GC "voice" is required.

7.1.2 Group Association Integrity Condition

The *group association integrity condition (IC)* specifies the set of conditions necessary to establish and maintain the group association. The integrity condition of a group association is composed of :

– the Active Group Integrity (AGI), that specifies conditions on the active group membership;

– the GA Topology Integrity (ATI), that specifies conditions on the group association topology.

When the integrity condition is not satisfied in the group association establishment phase, or is no more satisfied in the data transfer phase, one of the two policies defined here after can be applied :

– the association is released, that means the association does not exist any more. The policy is said a *hard IC management policy;*

– the association is suspended. In that case, data-transfers are not allowed until the integrity condition is satisfied again. The policy is said a *soft IC management policy*.

7.2 Group Connection

A *group connection (GC)* is the basic component of multi-peer communications. It can be considered as an actual data path for multicast data transmission. Three types of group connection are identified.

7.2.1 simplex group connection: A GC where one participant is the sender, in a multicast transmission mode, and the set of all other GC-participants are receivers.

7.2.2 duplex group connection: A GC where one participant is the sender, in a multicast transmission mode, and receiver of responses from the set of all other GC-participants. It is not possible to send/receive among the GC-participants.

7.2.3 N-plex group connection: A GC where any participant is a sender as well as a receiver in a multicast transmission mode, and if a sender sends a data unit all participants may receive it.

These three basic types of group connection defined here are thought to cover all the other types of connection as degenerate cases. For example, an *unicast simplex connection* is a degenerate case of the simplex group connection. An *unicast duplex connection* is a degenerate case of the duplex group connection or N-plex group connection. A $M \times N$ *group connection* of which M ($\leq$ N) members are intending to transmit multicast to the N members can be modeled as a degenerate case of the N-plex group connection. Even if an N-plex group connection may be seen as a collection of N of simplex group connection, modeling of an N-plex group connection is not equivalent to a set of N of simplex group connection at the service boundary. This would probably be a necessary technique in the protocol, but should not be permitted in the service primitives.

A simplex and duplex group connection has a central endpoint. The participant associated with this endpoint is called the *owner_of the group connection*. In case of N-plex group connection, a sender may be defined as the *owner* of the group connection. The three kinds of group connections are illustrated in Figure 4.

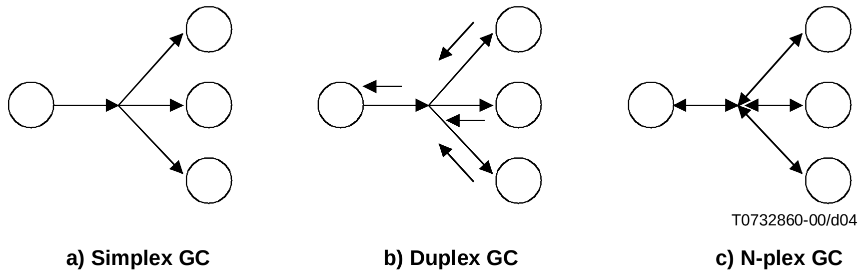

Figure 4/X.601 – Types of Group Connection

7.3 Relationship among GA, GC and Groups

Figures 5 and 6 show the relationships among different group types, GAs, and GCs.

Figure 5 shows the relationship of groups and nine group members A to I:

1) Multicast Group 1 = {A, B, C, D, F, G}, Multicast Group 2 = {A, B, C, E, H, I}

2) Registered Group 1 = {A, B, C, D, G}, Registered Group 2 = {A, B, C, E, H}

3) Enrolled Group 1 = {A, B, C, D}, Enrolled Group 2 = {A, B, C, E}

4) Active Group 1 = {A, C, D}, Active Group 2 = {A, B, C}

Figure 6 gives a typical example of a GA wherein two enrolled groups, one group association, and two group connections are defined as shown in Figure 5:

1) Group participants = { A, B, C, D}

2) Group Association = {GC1, GC2 }

3) GAEP_A = { GCEP_1, GCEP_2}

4) GAEP_ B = { GCEP_2}

5) GAEP_ C = { GCEP_1, GCEP_2 }

6) GAEP_ D = { GCEP_1}

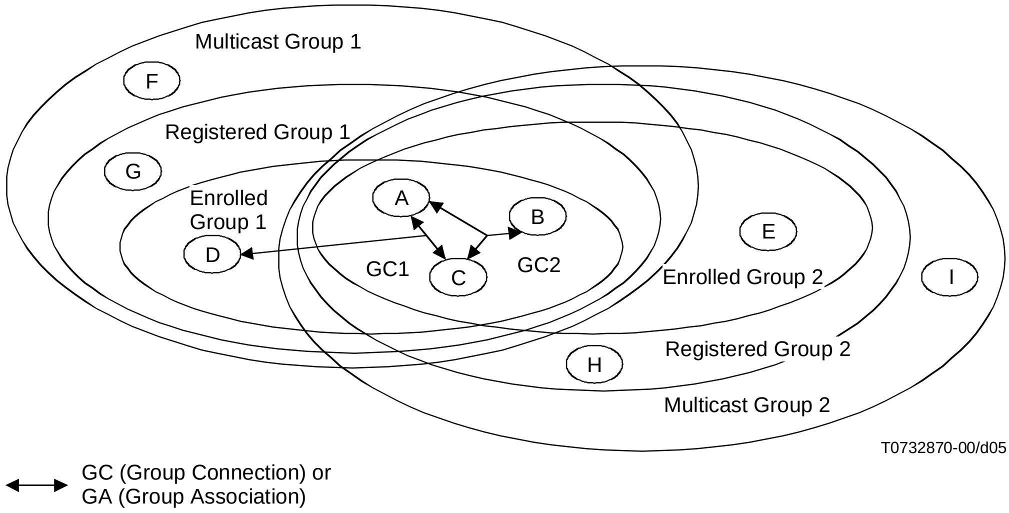

Figure 5/X.601 – Relationship among groups and group members

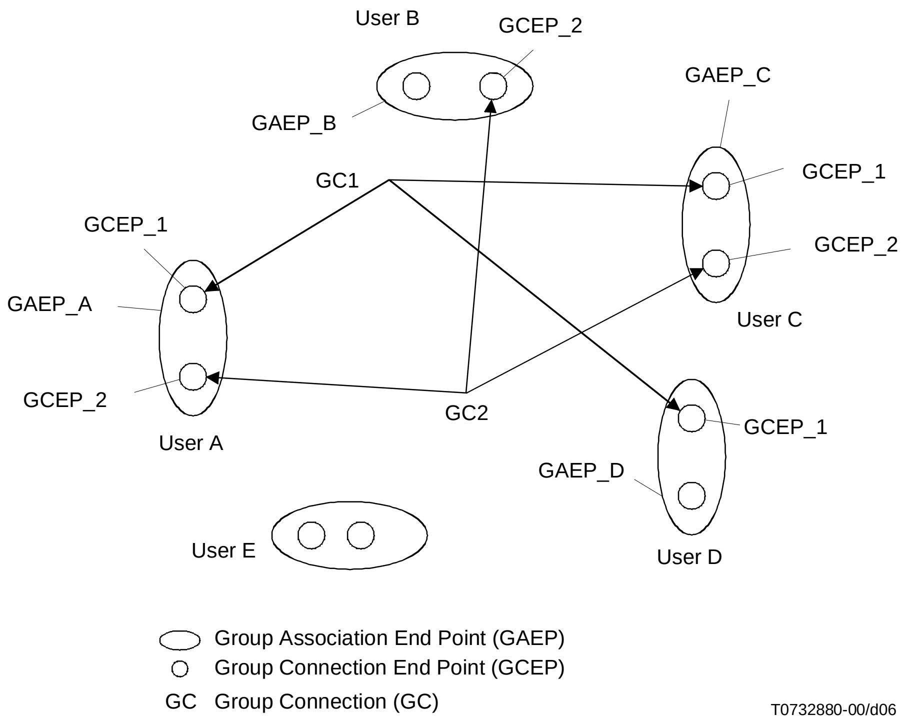

Figure 6/X.601 – Relationship between GA and GC

8 Quality of Service (QoS)

In the peer-to-peer case, one set of QoS values such as minimum value or peak value or average value is associated with each parameter on each direction of data transfer. In a multi-peer case, new dimensions are added to the problem due to the presence of several senders and receivers. In order to control any QoS in multi-peer communications, a set of QoS parameters is associated with each group connection. The QoS control is defined per group connection of a group-association. Within ITU-T and ISO/IEC JTC1, a QoS framework is developed that provides a common basis for the co-ordinated development and enhancement of a wide range of standards that specify or reference QoS requirements or mechanisms in an IT environment. Therefore, one can refer to this QoS framework when developing specifications dealing with QoS aspects.

8.1 Levels of QoS agreement

The term, level of agreement, is used to describe the agreed actions that will be taken by the service provider and/or the service users to maintain agreed levels of QoS. Three levels of agreement are defined: best efforts, compulsory, and guaranteed.

The weakest agreement is that all parties use their best efforts to meet the user requirements, but understand that there is no assurance that the agreed QoS will in fact be provided. In the compulsory level of agreement, the service must be aborted if the achieved QoS degrades below the agreed level; however, the agreed QoS level is not guaranteed, and indeed it may be deliberately degraded and the service aborted, for example to satisfy a higher-precedence demand for service. In the guaranteed level of agreement, the agreed QoS must be guaranteed so that the requested level will be met. This implies that the service will not be initiated unless it can be maintained within the specified limits.

8.2 QoS negotiation

In order to meet the needs of users and applications, quality of service should be agreed among the communicating entities before entering to the data transfer phase. This QoS agreement is made during connection establishment phase using the QoS negotiation procedure. Negotiation mechanisms are used to establish operating levels for QoS characteristics and to agree on the actions to be taken if these levels are not maintained.

For peer-to-peer communications, in the general case they require the participation of three parties (two service-users and a service-provider) and are therefore also described as three-party negotiation mechanisms. In the case of multi-peer communications, QoS characteristics need to be agreed on between the sender, the service-provider and multiple receivers.

Depending on the application, particular characteristics may be agreed on between the sender, the service-provider and each individual receiver independently or between the sender, the service-provider and all receivers jointly.

8.2.1 Receive diversity

From the viewpoint of receivers, a classification of QoS negotiation mechanisms is led to into those for "*connection-wide*", those for "*receiver-specific*" negotiation, and those for "*receiver-selected*" negotiation:

– *Connection-wide QoS negotiation:* negotiate the same value of a QoS characteristic for the sender, the service-provider and all receivers;

– *Receiver-specific QoS negotiation:* negotiate separate values of a QoS characteristic for each receiver, representing an agreement between the sender, the service-provider and that particular receiver.

– *Receiver-selected QoS negotiation:* negotiate different QoS values between sender and receiver. For example, different receivers may receive the data from the same sender at different QoS values not better than the transmit sending QoS. In this case, a translator may need to compensate differences between sender and receiver.

In the case of connectionless-mode service, negotiation is reduced to expressing the QoS requirements for individual data unit transfers, and this distinction disappears.

8.2.2 Transmit diversity

From the viewpoint of senders, a classification of QoS negotiation mechanisms is led to into those for "*homogeneous*" and those for "*heterogeneous*" negotiation:

– *homogeneous QoS negotiation:* negotiate to agree to a common set of transmit QoS values and for all senders transmit data at the same rate;

– *heterogeneous QoS negotiation:* negotiate for different senders to transmit data at different rates.

9 Phases of multi-peer communication

Multi-peer communication service may be achieved by seven distinct phases: registration, enrollment, activation, data transfer, deactivation, de-enrollment, de-registration.

9.1 Multicast-group creation phase

The first operation for the multi-peer communications is the creation of multicast-group. The multicast-group may be created by specifying the set of rules that define the future enrolled-groups membership. Before this operation, the multicast-group does not exist, and this operation will be performed by the group management facility.

As soon as an (N)-multicast group, i.e. possibly targeting (N)-enrolled-group, is reported to the group manager with (N)-group-membership criteria, an (N)-group-name is assigned to the group with an appropriate (N)-group-membership criteria. Then the information about this (N)-multicast group is to be known to the potential (N)-service users. Hence, an (N)-multicast-group can be identified by its (N)-group name in following actions by potential (N)-service users.

The criteria may be defined by means of a set of rules or by more arbitrary means such as a list. The group criteria may include group application name, group-name(s), proposed group-characteristics, and so on, depending on the hierarchy that the group is located. (N)-multicast-group is identified by the (N)-group name.

However, it may not yet be possible for any (N)-service user to determine the members of the (N)-multicast group, i.e. the rule defining the (N)-multicast group is known but it may not be feasible to determine all (N)-service users that satisfy the rule.

9.2 Registration/de registration phase

After creation of (N)-multicast-group, an (N)-service user should make it known to the group manager that he intends to be a member of (N)-enrolled-group.

Registration operation can be understood as simply letting the group manager know peer names and addresses of multicast group members. Thus (N)-registered-group is composed of members of (N)-multicast-group that have successfully performed the registration operation.

It may be possible for a certain (N)-service user, such as (N)-group owner, or (N)-group manager to determine the members of the (N)-registered group.

De-registration, the inverse operation of registration, is performed when an (N)-service user registration is deleted from the list of (N)-(N)-registered group members, or from the forwarding list in the intermediate systems.

9.3 Enrollment/de-enrollment phase

Enrollment operation allows an (N)-service user already registered to become a member of an (N)-enrolled-group. An (N)-enrolled-group is composed of members that have successfully performed the registration and the enrollment operations.

After creation of registered-group, an (N)-group address is assigned from the group manager, and is known to the registered group members. An (N)-enrolled group is actually in the position to be reached by means of at least one group-address or group-name. A member of an enrolled group can also be reached by its individual address. Then (N)-registered group members allow an intermediate system like a router to activate required transmission channel(s) and forward the data to itself.

The group enrollment operation accomplishes the whole "set-up" that is necessary to place an entity into the position to be instantiated as a member in an actual enrolled-group.

For example, in the Internet application, it is like selecting desirable group-connection(s) in a proposed group-association, letting those selected group-address(es) be activated in a designated router, which is capable of a multicast demon, using IGMP, and forwarding the data to the end-system. Even though the data is forwarded to the end system, it may not be received properly since the end system is not yet activated.

Also, it is like an SDR being advertised and each user selecting desirable group-connection(s) in a proposed group-association so that those selected group-address(es) can be activated in a designated router, which is capable of a multicast daemon, using IGMP, to forward the data to the end-system. And at the designated time, an end user activates the end system to receive data.

De-enrollment, the inverse operation of enrollment, is performed when an (N)-service user enrollment is deleted from the list of (N)-enrolled group members, or from the forwarding list in the intermediate systems. Then the (N)-enrolled group member permits an intermediate system like a router to release the transmission channel(s), so that the (N)-enrolled group member is in the position that can neither be reached by means of the (N)-group-address nor by the (N)-group-name.

The state transition diagram, presented in Figure 7 describes different states of a group member. This diagram does not take into account the interactions between different layers. The initial state of a member is undefined state, where no criteria for establishing a group are defined. When a certain criteria for a group is defined where a multicast group is defined, then a group member goes into the defined state. When a group member of a multicast group has passed the registration operation successfully, then the group member goes into the registered state. When he has further passed the enrollment operation successfully, he then goes into the enrolled state where the enrolled group is established. Finally, if a member of an enrolled group has passed the activation operation successfully, then he goes into the activated state where the active group is established.

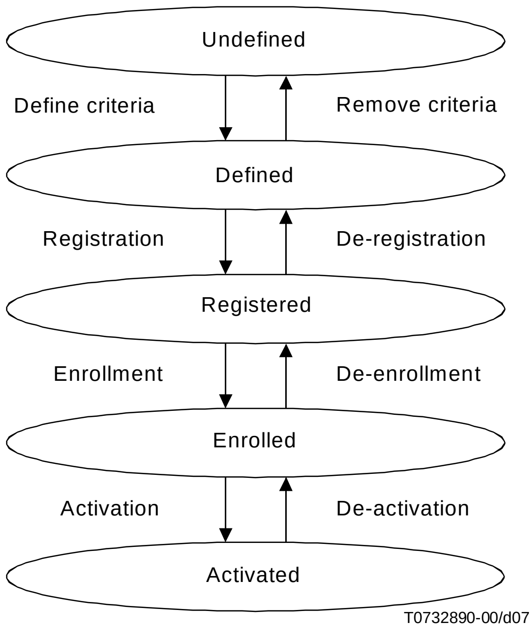

Figure 7/X.601 – State transition diagram of a group member

9.4 Activation/deactivation

The activation phase establishes the shared state of a specific instance of a multi-peer communication based on information established during the enrollment phase among the participants. The successful completion of the activation phase leaves the entity in the data transfer phase. An (N)-active-group is composed of (N)-enrolled group members that have completed the activation phase.

It is during the activation phase that the specific QoS requirements acceptable for data transfer must be made or modified if they were not made during the enrollment phase.

In this phase, group association and/or group connection is created. The member that initiates the creation of GA and/or GC is called the initiator.

According to the (N)-group membership characteristics, any (N)-group member may join the existing (N)-group association and/or (N)-group connection that has already been created. Two kinds of join operation can be distinguished:

– *Join-Invitation* operation : allows existing participants to invite an entity to join GA or GC;

– *Join-Calling* operation : allows an entity to request to participate in an existing GA or GC.

This operation assumes that the (N)-group member knows that the (N)-GA or (N)-GC have been previously established, and knows how to address the (N)-GA or (N)-GC. The possibility of joining an existing (N)-GA or (N)-GC depends on the (N)-group membership characteristics. The join operation is allowed only for a dynamic group. In case of a closed group, the entity must belong to the (N)-enrolled-group. In case of an open group, the entity may not belong to the (N)-enrolled-group.

During the activation phase, some (N)-enrolled group members may accept the establishment of the (N)-GA and/or (N)-GC while others may reject it. For the (N)-service-provider to be able to decide if the establishment is successful or not, a condition needs to be verified. Such a condition, called the establishment condition, expresses constraints such as AGI conditions and QoS requirements. At the end of the establishment operation, if this establishment condition is verified, a new GA and/or GC is created and/or a group member is joined to an existing GA and/or GC.

The reverse of a join operation is the leave operation. This operation allows an (N)-group participant to leave an (N)-GA and/or (N)-GC. This operation may be initiated by the participant that leaves the (N)-GA and/or (N)-GC or by another (N)-group participant. The leave operation may also be initiated by the (N)-service-provider when it is no longer in a situation to satisfy the requested service with the requested QoS. Depending on the integrity condition, this operation may lead to the release of the (N)-GA and/or (N)-GC.

The activation and its reverse operation, deactivation, are the traditional operations of putting a facility "on-line" and taking it "off-line" without deleting the system's knowledge that the facility exists.

The deactivation operation is performed to terminate an (N)-GA and/or (N)-GC. The deactivation may occur at any time and may or may not be abrupt. It may be initiated either by the (N)-service-user or the (N)-service-provider, for instance, when the integrity condition is no longer satisfied in the case of hard integrity condition management policy. It is destructive which means that, as a result of the release, the (N)-GA and/or (N)-GC does not exist any longer. The completion of the deactivation phase returns the entity to the enrollment phase.

9.5 Data transfer phase

The data transfer phase involves the actual transfer of data among the (N)-group participants who belong to the (N)-active group. It may include functions to maintain the QoS, either in terms of error control or resource reservation. The data transfers are allowed among the participants only if the integrity condition is satisfied. If the integrity condition is no longer satisfied during the data transfer, the data transfer operation may be suspended (soft) or terminated (hard) according to the integrity condition management policy.