ANCP
Internet-Draft
Intended status: Standards Track
Expires: August 13, 2011

F. Le Faucheur
Cisco
R. Maglione
Telecom Italia
T. Taylor
Huawei

February 9, 2011

Multicast Control Extensions for ANCP
draft-ietf-ancp-mc-extensions-04.txt

Abstract

This document specifies the extensions to the Access Node Control Protocol required for support of the multicast use cases defined in the Access Node Control Protocol framework document and one additional use case described in this document. These use cases are organized into the following ANCP capabilities:

- NAS-initiated multicast replication;
- conditional access with white and black lists;
- conditional access with grey lists;
- bandwidth delegation;
- committed bandwidth reporting.

These capabilities may be combined according to the rules given in this specification.

This Internet-Draft will expire on August 13, 2011.

Table of Contents

1. Introduction

[RFC5851] defines a framework and requirements for an Access Node control mechanism between a Network Access Server (NAS) and an Access Node (e.g. a Digital Subscriber Line Access Multiplexer (DSLAM)) in a multi-service reference architecture in order to perform QoS-related, service-related and subscriber-related operations. [I-D.ietf-ancp-protocol] specifies a protocol for Access Node Control in broadband networks in line with this framework.

[I-D.ietf-ancp-protocol] supports three use cases defined in [RFC5851], specifically for DSL access: the DSL Topology Discovery use case, the DSL Line Configuration use case and the DSL Remote Connectivity Test use case. However, it does not support the multicast use cases defined in [RFC5851]. The present document specifies the extensions to the Access Node Control Protocol required for support of these multicast use cases. In addition, it supports the Committed Bandwidth Reporting use case, described below. In terms of the ANCP protocol, these use cases are organized into five capabilities:

- NAS-initiated multicast replication;
- conditional access with white and black lists;
- conditional access with grey lists;
- bandwidth delegation;
- committed bandwidth reporting.

NAS-initiated multicast replication assumes that multicast "join" and "leave" requests are terminated on the NAS, or that the NAS receives requests to establish multicast sessions through other means (e.g., application-level signalling). The NAS sends commands to the AN to start or stop replication of specific multicast flows on specific subscriber ports. This use case is described briefly in the next-to-last paragraph of Section 3.4 of [RFC5851].

Conditional access is described in Sections 3.4.1 and 3.4.2.3 of [RFC5851], with the latter section particularly applicable to operation with white and black lists only. In case of "conditional access with white and black lists", multicast join and leave requests are terminated at the AN and accepted or ignored in accordance with the direction provided by white and black lists respectively. The white and black lists are provisioned per port at startup time and may be modified thereafter. The NAS may enable admission control of white-listed flows by appropriate provisioning.

Conditional access with grey lists is similar to conditional access with white lists, except that before accepting any request matching a grey list entry, the AN sends a request to the NAS for permission to replicate the flow. Again, the NAS can enable admission control of grey-listed flows at the AN.

Bandwidth delegation is described in Section 3.4.2.1 of [RFC5851]. It allows flexible sharing of total video bandwidth on an access line between the AN and the NAS. One application of such bandwidth sharing is where the AN does multicast admission control, while the NAS or Policy Server does unicast admission control. In that case, bandwidth delegation allows dynamic sharing of bandwidth between unicast and multicast video traffic on each access line.

Committed bandwidth reporting is described below, in Section 3.4. The AN reports the amount of multicast bandwidth it has granted to a given access line each time that value changes. These reports may be buffered for a NAS-provisionable interval so that reports for multiple access lines can be bundled into the same message.

The formal specification of the behaviours associated with each of these capabilities, singly and in combination, is given in Section 6.

In addition to the multicast service processing behaviour just sketched, the definition of each capability includes support for the multicast accounting and reporting services described in Section 3.4.3 of [RFC5851]. Because of this common content and because of other protocol overlaps between the different capabilities, the protocol descriptions for the multicast extensions specified in this document are merged into a single non-redundant narrative. Tables in Section 6 then indicate the specific sub-sections of the protocol description that have to be implemented to support each capability.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

The expression "delegated bandwidth" is used as a shorter way of saying: "the total amount of video bandwidth delegated to the AN for multicast admission control".

3. Multicast Use Cases

Quoting from [RFC5851]:

"... the Access Node, aggregation node(s) and the NAS must all be involved in the multicast replication process. This avoids that several copies of the same stream are sent within the access/aggregation network. In case of an Ethernet-based access/aggregation network, this may, for example, be achieved by means of IGMP snooping or IGMP proxy in the Access Node and aggregation node(s). By introducing IGMP processing in the access/aggregation nodes, the multicast replication process is now divided between the NAS, the aggregation node(s) and Access Nodes. In order to ensure backward compatibility with the ATM-based model, the NAS, aggregation node and Access Node need to behave as a single logical device. This logical device must have exactly the same functionality as the NAS in the ATM access/aggregation network. The Access Node Control Mechanism can be used to make sure that this logical/functional equivalence is achieved by exchanging the necessary information between the Access Node and the NAS. "

[RFC5851] describes the use cases for ANCP associated with such multicast operations, and identifies the associated ANCP requirements. The present section describes a subset of these use cases as background to facilitate reading of this document, but the reader is refered to [RFC5851] for a more exhaustive description of the ANCP multicast use cases. Detailed example message flows can also be found in Appendix A.

3.1. NAS Initiated Multicast Replication Control Use Case

3.1.1. Goals

One option for multicast handling is for the subscriber to communicate the "join/leave" information to the NAS. This can be done for instance by terminating all subscriber IGMP ([RFC3376]) or MLD ([RFC2710], [RFC3810]) signaling on the NAS. Another example could be a subscriber using some form of application level signaling, which is redirected to the NAS. In any case, this option is transparent to the access and aggregation network. In this scenario, the NAS uses ANCP to create and remove replication state in the AN for efficient multicast replication. Thus, the NAS only sends a single copy of the multicast stream towards the AN, which in turn performs replication to multiple subscribers as instructed by the NAS via ANCP. The NAS first performs conditional access and multicast admission control when processing multicast join requests, and only creates replication state in the AN if admission succeeds.

3.1.2. Message Flow

With the NAS-initiated use case, a Multicast Replication Control Message is sent by the NAS to the AN with a directive to either join or leave one (or more) multicast flow(s). In the example message flow, the AN uses a Generic Response message to convey the outcome of the directive. Figure 1 illustrates such an ANCP message exchange as well as the associated AN behavior.

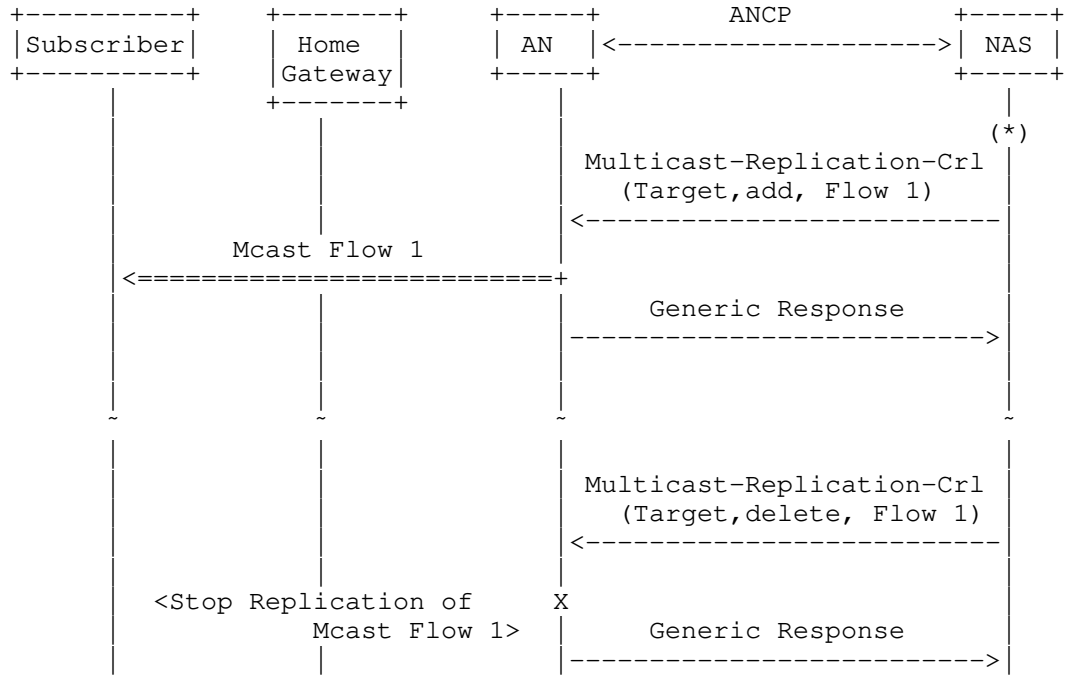

(*) The NAS may optionally seek direction from an external Authorization/Policy Server before admitting the flow.

Figure 1: NAS Initiated Multicast Replication Control

3.2. Conditional Access and Admission Control Use Case

3.2.1. Goals

One option for multicast handling is for the access/aggregation nodes to participate in IGMP/MLD processing (e.g. via IGMP/MLD snooping). In this scenario, on detecting a join/leave request from an enduser for a multicast flow (in the grey list), the AN uses ANCP to request conditional access and admission control decision from the NAS. In

turn, after conditional access and admission control checks, the NAS uses ANCP to instruct the AN to change the replication states accordingly.

3.2.2. Message Flow

For support of the conditional access and admission control use case, on detection of an IGMP/MLD Join, the AN sends an Admission Control message to the NAS to request conditional access and admission control check. In case of positive outcome, the NAS sends a Multicast Replication Control Message to the AN with a directive to replicate the multicast flow to the corresponding user. Similarly on detection of an IGMP/MLD leave, an Admission Control message is sent by the AN to the NAS to keep the NAS aware of user departure for the flow. This message flow is illustrated in Figure 2.

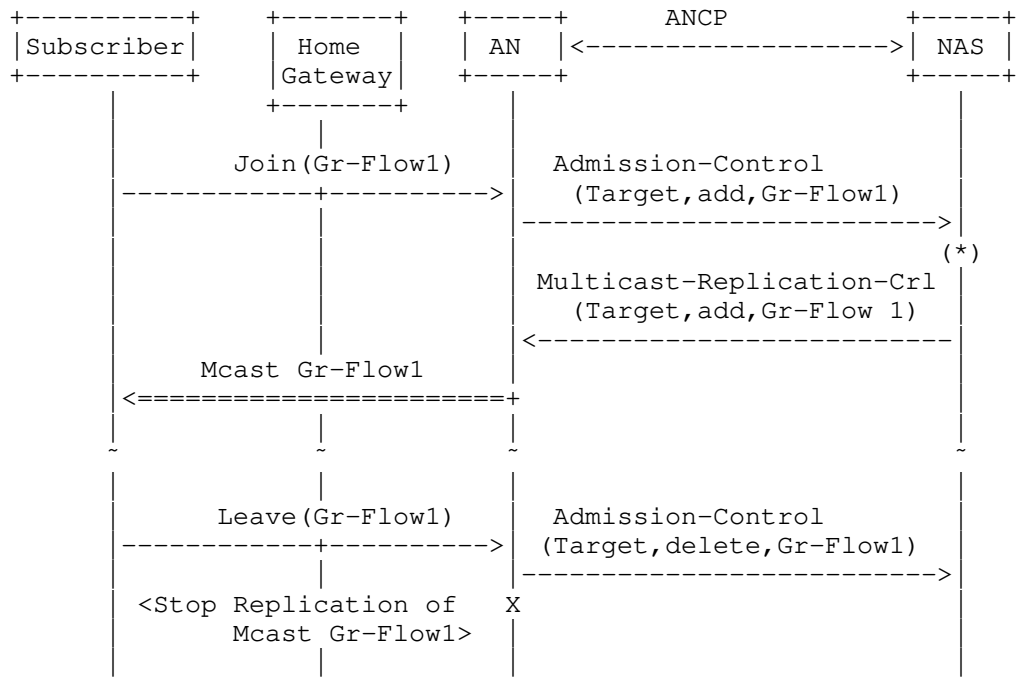

Gr-Flow1: a multicast flow matching the grey list for that port

(*) The NAS may optionally seek direction from an external Authorization/Policy Server before admitting the flow.

Figure 2: Multicast Conditional Access and Admission Control

3.3. Multicast Flow Reporting Use Case

3.3.1. Goals

The Multicast flow reporting use case allows the NAS to asynchronously query the AN to obtain an instantaneous status report related to multicast flows currently replicated by the AN.

3.3.2. Message Flow

The NAS sends a Multicast Flow Query Request message to the AN in order to query the AN about information such as which multicast flows are currently active on a given AN port or which ports are currently replicating a given multicast flow. The AN conveys the requested information to the NAS in a Multicast Flow Query Response message. This message flow is illustrated in Figure 3.

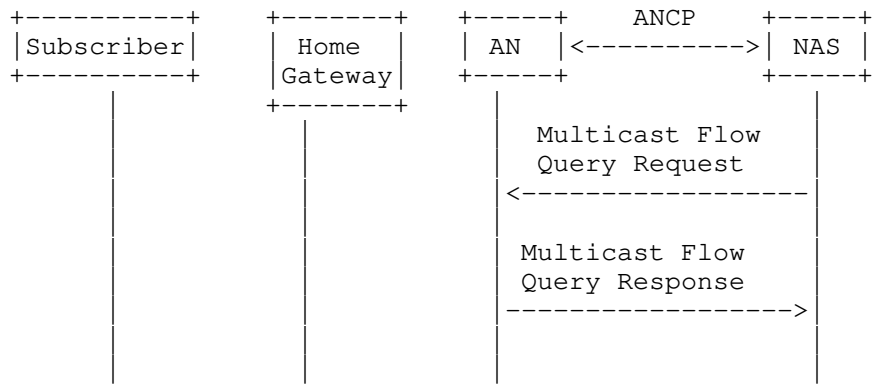

Figure 3: Multicast Flow Reporting

3.4. Committed Bandwidth Reporting Use Case

3.4.1. Goals

The committed bandwidth reporting use case allows the NAS to maintain current awareness of how much multicast bandwidth the AN has committed to a given access line, so that the NAS can adjust its forwarding scheduler to ensure the associated QoS. Note that this involves a finer level of detail than provided by bandwidth delegation, since the amount of delegated bandwidth is an upper limit on the amount of bandwidth committed rather than an actual value. To reduce the volume of messaging, reports from the AN may be buffered so that one message reports on changes for multiple access lines.

3.4.2. Message Flow

The message flow associated with this use case is shown in Figure 4. The figure assumes that a non-zero buffering interval was previously provisioned on the AN.

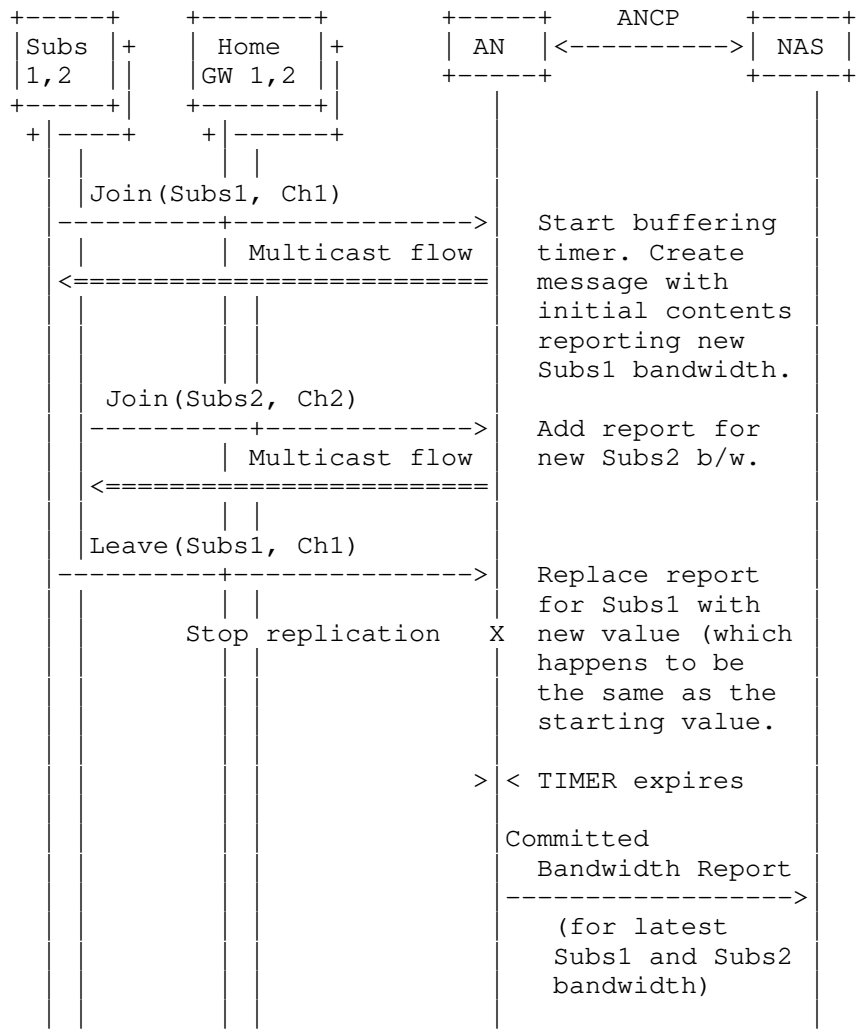

Figure 4: Message Flow For Committed Bandwidth Reporting

4. ANCP Messages

This section defines new ANCP messages and new usage of existing ANCP messages as well as procedures associated with the use of these messages.

4.1. Provisioning Message

Section 6.1.2 of [I-D.ietf-ancp-protocol] defines the Provisioning message that is sent by the NAS to the AN to provision information in the AN.

The present document specifies that the Provisioning message MAY be used by the NAS to provision multicast-related information (e.g. multicast service profiles). The ANCP Provisioning message payload MAY contain:

- one or more instances of the Multicast-Service-Profile TLV. The Multicast- Service-Profile TLV is defined in the present document in Section 5.1. Each instance of the Multicast-Service-Profile TLV contains a multicast service profile name and one or more list actions. A list action consists of an action (add, delete, replace), a list type (White, Black, or Grey), and list content (multicast source and group addresses).

- an instance of the White-List-CAC TLV. The White-List-CAC TLV is defined in Section 5.6. If present, this TLV indicates that the AN is required to do admission control before replicating White-listed flows.

- an instance of the MRepCtl-CAC TLV. The MRepCtl-CAC TLV is defined in Section 5.7. If present, this TLV indicates that the AN is required to do admission control before replicating flows specified in Multicast Replication Control messages.

- an instance of the Report-Buffering-Time TLV. The Report-Buffering-Time TLV is defined in Section 5.12. If present, this TLV indicates Committed Bandwidth Report messages should be buffered for the amount of time given by the TLV before being transmitted to the NAS.

See Section 6 for information on which multicast capabilities require support of these TLVs in the Provisioning message.

4.1.1. Sender Behaviour

When directed by the Policy Server or by management action, the NAS sends the Provisioning message to initially provision or to update

the White, Black, and/or Grey multicast channel lists associated with a set of named multicast service profiles, or to enable the AN to perform admission control for specific classes of flows.

To provision or update a multicast service profile, the NAS MUST include within the message one or more instances of the Multicast- Service- Profile TLV specifying the content to be provisioned or updated. The NAS SHOULD NOT include any list type (White, Black, or Grey) that is not supported by the set of multicast capabilities negotiated between the NAS and the AN. The NAS MUST NOT use the Provisioning message to send instances of the Multicast- Service- Profile TLV to the AN unless the Multicast-Service-Profile TLV is supported by the set of multicast capabilities negotiated between the NAS and the AN.

To require admission control to be performed at the AN on White-listed flows, the NAS MUST include a copy of the White-List-CAC TLV in the Provisioning message. The White-List-CAC TLV MUST NOT be provided unless the negotiated set of capabilities includes conditional access with White and Black lists.

To require admission control to be performed at the AN on Grey-listed flows or on NAS-initiated flows, the NAS MUST include a copy of the MRepCtl-CAC TLV in the Provisioning message. The MRepCtl-CAC TLV MUST NOT be provided unless the negotiated set of capabilities includes NAS-initiated replication control or conditional access with Grey lists.

To require buffering of Committed Bandwidth Report messages so that reports for multiple access lines can be included in the same message, the NAS MUST include a copy of the Report-Buffering-Time TLV containing a non-zero time value in a Provisioning message sent to the AN. The Report-Buffering-Time TLV MUST NOT be provided unless the negotiated set of capabilities includes committed bandwidth reporting.

4.1.2. Receiver Behaviour

The receiving AN provisions/updates the White, Black, and/or Grey lists associated with the multicast service profile names contained in the Multicast-Service-Profile TLV instances within the message according to the contents of the associated List-Action TLVs. The AN MUST process List-Action TLVs in the order in which they appear within the message. The AN MUST ignore instances of the List-Action TLV referring to any list type (White, Black, or Grey) that is not supported by the set of multicast capabilities negotiated between the NAS and the AN.

When a new multicast service profile is identified by a Multicast-Service-Profile TLV, the initial state of all lists associated with that profile according to the negotiated set of multicast capabilities is empty until changed by the contents of Multicast-Service-Profile TLVs.

The receipt of a Provisioning message containing updates to an existing multicast service profile subsequent to startup will cause the AN to review the status of active flows on all ports to which that profile has been assigned. For further details, see Section 6.

If the White-List-CAC and/or MRepCtl-CAC TLV is present in the Provisioning message and the respective associated capabilities have been negotiated, the AN prepares (or continues) to do connection admission control on the indicated class(es) of flow. If one or both of these TLVs was present in an earlier Provisioning message but is absent in the latest message received, the AN ceases to do connection admission control on the indicated class(es) of flow.

The buffering time specified in an instance of the Report-Buffering-Time TLV applies to only to Committed Bandwidth Report messages initiated after the new buffering time is received at the AN, not to any message already in the process of accumulation.

As indicated in [I-D.ietf-ancp-protocol], the AN MUST NOT reply to the Provisioning message if it processed it successfully. If an error prevents successful processing of the message content, the AN MUST return a Generic Response message as defined in [I-D.ietf-ancp-protocol], containing a Status-Info TLV with the appropriate content describing the error. For this purpose, the presence of a list type in a Multicast-Service-Profile TLV which was ignored because it was not supported by the negotiated set of capabilities is not considered to be an error.

4.2. Port Management Message

As specified in [I-D.ietf-ancp-protocol], the NAS may send DSL line configuration information to the AN ("ANCP based DSL Line Configuration" use case) using GSMP Port Management messages modified to contain additional information. See Section 5.3.3 of [I-D.ietf-ancp-protocol] for details.

This document specifies that the Port Management message MAY also include either or both of the following TLVs:

- Multicast-Service-Profile-Name TLV (defined in Section 5.2). This TLV associates a Multicast Service Profile with the Access Port specified by the extension block.

- o Bandwidth-Allocation TLV (defined in Section 5.5). This TLV specifies the total multicast bandwidth available to the AN for admission control at the Access Port.

4.2.1. Sender Behaviour

The NAS sends the Port Management message at startup time to initialize parameters associated with the Access Port specified in the message and with the multicast capabilities negotiated between the NAS and the AN. The NAS MAY send additional Port Management messages subsequent to startup, to update or, in the case of the Bandwidth-Allocation TLV, reset these parameters. If the NAS includes a Multicast-Service-Profile-Name TLV in the Port Management message, the name MUST match a profile name provided in a Multicast- Service- Profile TLV in a prior Provisioning message. The NAS MUST NOT include a TLV unless it is supported by the set of multicast capabilities negotiated between the NAS and the AN. See Section 6 for further information.

4.2.2. Receiver Behaviour

If the Port Management message contains a Multicast-Service-Profile-Name TLV, the AN associates the named profile with the specified Access Port. This association replaces any previous association. That is, a given Access Port is associated with at most one multicast service profile. The replacement of one multicast service profile with another will cause the AN to review the status of all active flows on the target port. For further details see Section 6.

If the Port Management message contains a Bandwidth-Allocation TLV, the AN adopts this as the current value of its total multicast bandwidth limit for the target port. If the AN has already committed multicast bandwidth exceeding the amount given in the Bandwidth-Allocation TLV, the AN SHOULD NOT discontinue any multicast streams in order to bring bandwidth down to within the new limit. However, the AN MUST NOT admit new multicast streams that are subject to admission control until it can do so within the limit specified by the Bandwidth-Allocation TLV.

If the Port Management request cannot be processed due to error and the Result field of the request is Nack (0x1) or AckAll (0x2), the AN SHOULD add a Status-Info TLV to the Extension Value field in its reply if this will provide useful information beyond what is provided by the Code value returned in the response header. In particular, if the name within the Multicast-Service-Profile-Name TLV does not match a profile name given in a prior Provisioning message, the AN SHOULD return a reply where the Code field in the header indicates "Invalid TLV value" (85), the Error Message field in the Status-Info TLV

contains the text "Multicast profile name not provisioned", and the Status-Info TLV contains a copy of the Multicast-Service-Profile-Name TLV.

4.3. Multicast Replication Control Message

This section defines a new message called the Multicast Replication Control message. The Multicast Replication Control message is sent by the NAS to the AN with one or more directives to add (join) or delete (leave) a multicast flow on a target object identified in the content of the message.

The Message Type for the Multicast Replication Control message is 144.

The ANCP Multicast Replication Control message payload contains the following TLVs:

- Target TLV: The Target TLV is defined in [I-D.ietf-ancp-protocol]. It MUST appear once and only once. It is encoded as specified in [I-D.ietf-ancp-protocol] or extensions and identifies the AN port subject to the request for admission or release.

- Command TLV: The Command TLV is defined in [I-D.ietf-ancp-protocol]. It MUST be present. It MAY appear multiple times.

As [I-D.ietf-ancp-protocol] indicates, the contents of the Command Info field within the Command TLV are specific to the message in which the TLV occurs. For the Multicast Replication Control Message, these contents consist of:

- a Command Code field;

- an Accounting field;

- an instance of the Multicast-Flow TLV.

Figure 5 illustrates the complete Command TLV with the contents specific to the Multicast Replication Control message.

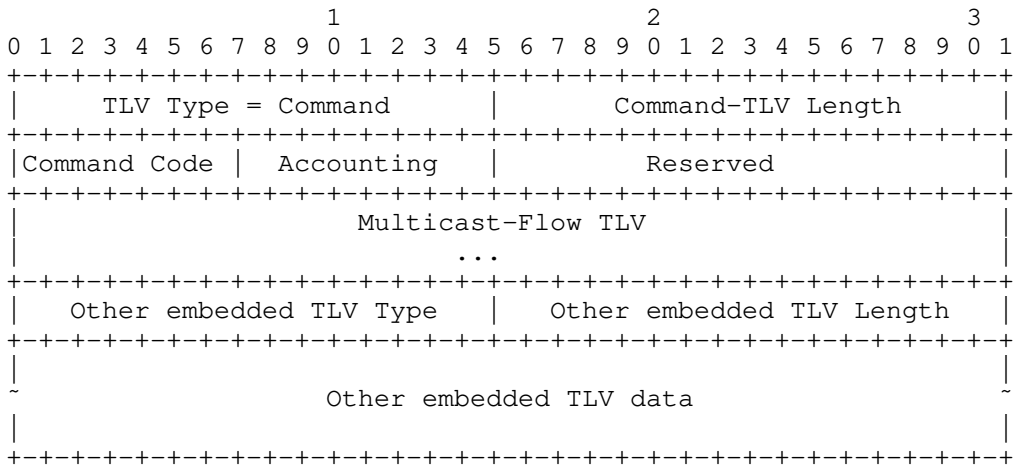

Figure 5: Contents of the Command TLV in the Multicast Replication Control Message

Command Code:

Command directive:

0x01 - Add;

0x02 - Delete;

0x03 - Delete All;

0x04 - Admission Control Reject;

0x05 - Conditional Access Reject;

0x06 - Admission Control and Conditional Access Reject.

Directives 0x04 through 0x06 are used as described in Section 4.4.2

Accounting:

Meaningful only when the Command Code is "Add" (0x01). In that case, 0x00 indicates no flow accounting, 0x01 indicates that octet accounting for the flow is to commence. The Accounting field MUST be set to 0x00 for other Command Code values.

Reserved:
: Reserved for future use. MUST be set to 0x0000 by the sender and ignored by the receiver.

Multicast-Flow TLV:
: An instance of the Multicast-Flow TLV (Section 5.11) specifying the flow to be added or deleted. The Multicast-Flow TLV MUST be omitted if the Command Code has value "Delete All" (0x03).

Other embedded TLV:
: No other embedded TLVs are currently specified within the Multicast Replication Control message/Command TLV. Unrecognized embedded TLVs SHOULD be silently discarded.

The figure below is an example of a Multicast Replication Control message that would result in a swap from multicast SSM flows 192.0.2.1, 233.252.0.2, to 192.0.2.2, 233.252.0.3 on the Target identified by the "Access Loop Circuit ID":

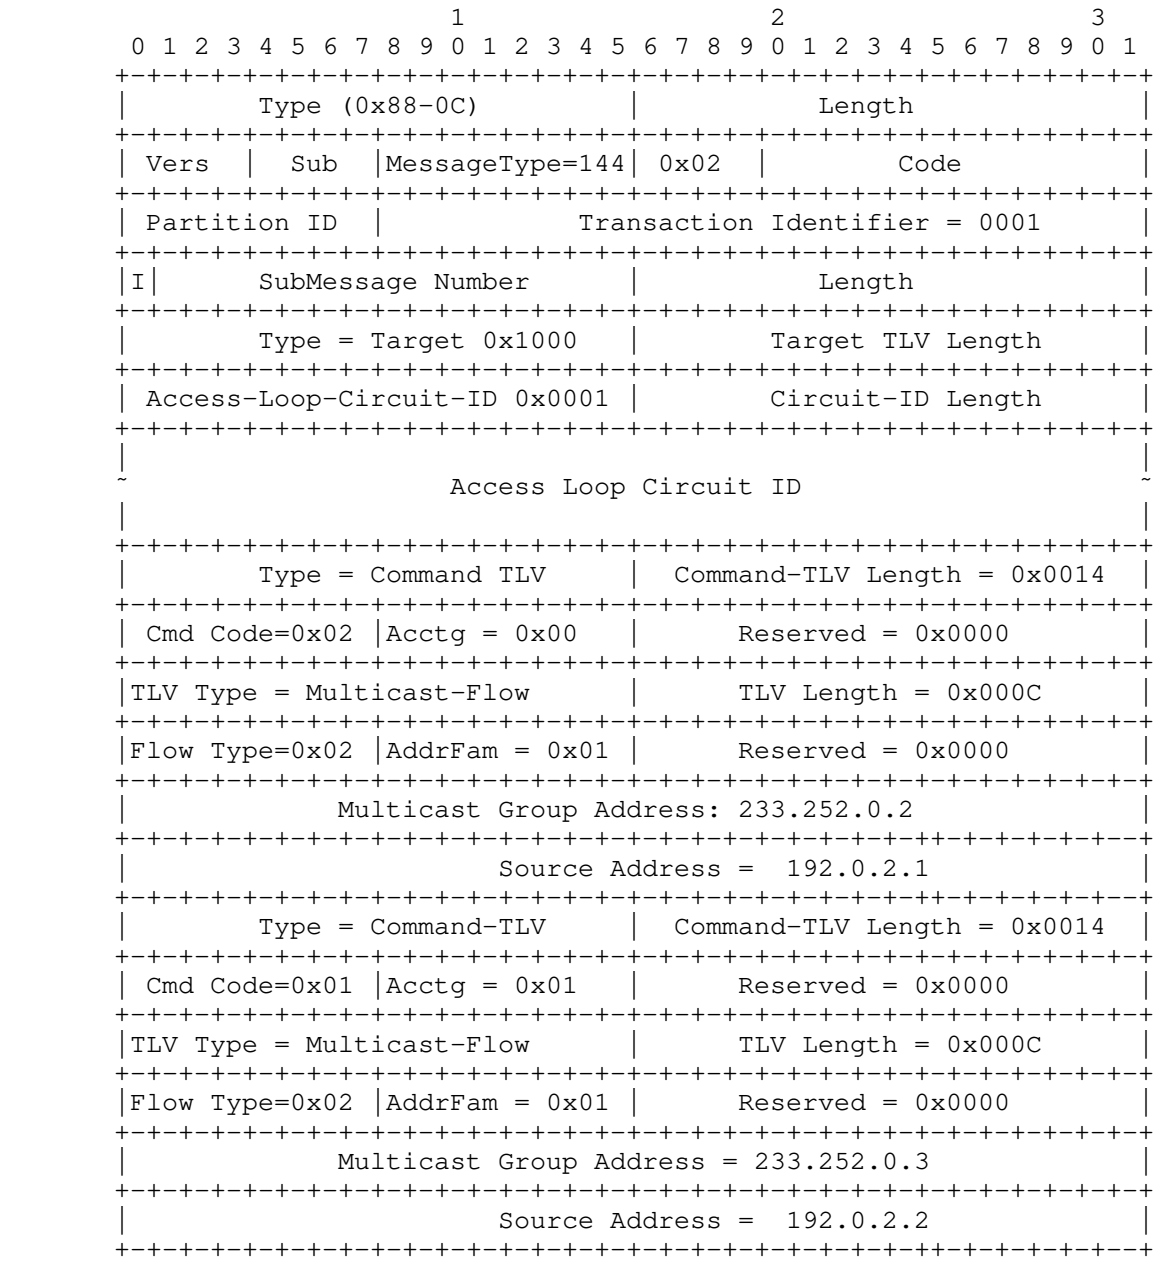
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
Type (0x88-0C)
Length
Vers
Sub
MessageType=144
0x02
Code
Partition ID
Transaction Identifier = 0001
I
SubMessage Number
Length
Type = Target 0x1000
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID
Type = Command TLV
Command-TLV Length = 0x0014
Cmd Code=0x02
Acctg = 0x00
Reserved = 0x0000
TLV Type = Multicast-Flow
TLV Length = 0x000C
Flow Type=0x02
AddrFam = 0x01
Reserved = 0x0000
Multicast Group Address: 233.252.0.2
Source Address = 192.0.2.1
Type = Command-TLV
Command-TLV Length = 0x0014
Cmd Code=0x01
Acctg = 0x01
Reserved = 0x0000
TLV Type = Multicast-Flow
TLV Length = 0x000C
Flow Type=0x02
AddrFam = 0x01
Reserved = 0x0000
Multicast Group Address = 233.252.0.3
Source Address = 192.0.2.2

4.3.1. Sender Behaviour

The NAS MAY issue a Multicast Replication Control message to the AN to convey one or more directives to add (join) or delete (leave) one or more multicast flows.

The NAS MAY send this message on its own initiative to support the NAS initiated Multicast Control use case presented in [RFC5851] and summarized in Section 3.1. In that case, the NAS MUST set the Result field to AckAll (0x2) or Nack (0x1) according to its requirements.

The NAS MAY also send this message in response to a Multicast Admission Control message (defined in Section 4.4) received from the AN to support the conditional access and admission control use case presented in [RFC5851] and summarized in Section 3.2. In that case, the NAS MUST set the Result field to NAck (0x1).

In either case, the sender MUST populate the Code field with the value 0x000 and the ANCP Transaction Identifier field with a unique value, as described in Section 4.4.1 of [I-D.ietf-ancp-protocol].

Each Multicast Replication Control Message MUST contain one or more commands, each encapsulated in its own Command TLV. The sender MUST use a separate Command TLV for each distinct multicast flow.

When the order of processing of two commands does not matter, the commands MUST be transmitted in separate Multicast Replication Control messages.

4.3.2. Receiver Behaviour

When successive commands (in the same or different messages) relate to the same Target and multicast flow, the state of each feature controlled or affected by attributes received in the Multicast Replication Control message, SHALL be as set by the last command or message referring to that target and flow and containing the controlling attribute. As an example, successive Multicast Replication Control messages containing add commands for a given port and flow, but differing only in the Accounting field setting SHALL be interpreted to mean that the state of the accounting feature is as set in the final command received, but all other features are as set in the initial message.

If more than one Command TLV is present in a Multicast Replication Control message, the AN MUST act on the commands in the order in which they are presented in the message. The AN SHALL assign a sequence number to each command in a given Multicast Replication Control message, starting from 0x01 for the first command.

If a Command TLV adds a flow and the AN is performing admission control for Multicast Replication Control messages, then the AN MUST perform admission control before replicating the flow. If the admission control check fails, the AN MUST treat the failure as an error as described below. The appropriate Code value for the response is 18 (0x012) "Insufficient resources".

If the AN processes the complete Multicast Replication Control message successfully and the Result field of the Multicast Replication Control message was set to AckAll (0x2), the AN MUST respond with a Generic Response message where the Result field is set to Success (0x3), the Code field is set to 0x000, and the Transaction Identifier field is copied from the Multicast Replication Control message. The body of the response MAY be empty or MAY be copied from the Multicast Replication Control message.

If the AN processes the complete Multicast Replication Control message successfully and the Result field of the Multicast Replication Control message was set to Nack (0x1), the AN MUST NOT respond to the message.

The processing/execution of multiple commands contained in a single Multicast Control message MUST be interrupted at the first error encountered, and the remaining commands in the Multicast Replication Control message discarded.

If the AN detects an error in a received Multicast Replication Control message and the Result field in that message was set to Nack (0x1) or AckAll(0x2), the AN MUST generate a Generic Response message providing error information to the NAS. This specification identifies the following new Code values beyond those specified in [I-D.ietf-ancp-protocol], which MAY be used in a Generic Response sent in reply to a Multicast Replication Control message:

100 Command error. This SHOULD be reported for the case that an invalid command code has been received.

101 Bad flow address. This SHOULD be reported for the following cases:

* unsupported address family;

* source address present for an ASM flow, or absent for an SSM flow.

102 Multicast flow does not exist. This SHOULD be reported if the NAS attempts to delete a flow that is not enabled.

A Generic Response message responding to the Multicast Replication Control message and containing one of the above Code values MUST include a Status-Info TLV which includes one or two embedded TLVs as follows:

- a Sequence-Number TLV as described in Section 5.4, giving the sequence number of the failed command, MUST be included;

- the failed Command TLV itself SHOULD be included.

  Note that the Error Message field of the Status-Info TLV MAY be used to report more details than implied by the Code value in the message header. For example, the Code value could be 101 and the Error Message field could contain the text: "Source address present for ASM flow".

4.4. Multicast Admission Control Message

This section defines a new message called the Multicast Admission Control message. The Multicast Admission Control message is sent by the AN to the NAS to request admission of a multicast flow, or to notify of the removal of a multicast flow, for a given target.

The Message Type for the Multicast Admission Control message is 145.

The ANCP Multicast Admission Control message payload contains two TLVs:

- Target TLV: The Target TLV is defined in [I-D.ietf-ancp-protocol]. It MUST appear once and only once in the Multicast Admission Control message. It is encoded as specified in [I-D.ietf-ancp-protocol] or extensions and identifies the AN port subject to the request for admission or release.

- Command TLV: The Command TLV is defined in [I-D.ietf-ancp-protocol]. It MUST be present. If it appears more than once, only the first instance is considered meaningful in the present version of this specification and the other instances are ignored.

Informative note:

  In the future, the specification of the Admission Control message may be extended to allow transport of more than a single directive (e.g. to carry both a leave from one group and a join to another

group for the same Target). It is expected that this would support a similar notion of strict sequenced processing as currently defined for handling multiple directives in the Multicast Replication Control message whereby all directives following the first directive that can not be executed are not executed either. When the strict sequenced processing of the directives is not required the directives are distributed across separate messages.

The Command TLV has the same contents as were described above for the Multicast Replication Control message, with the following additions:

- a Request-Source-IP TLV MAY be appended to the Command TLV as an additional embedded TLV;
- similarly, a Request-Source-MAC TLV MAY be appended to the Command TLV as an additional embedded TLV.

Note that the Command TLV length includes the length of any embedded TLVs, including the embedded TLV headers.

4.4.1. Sender Behaviour

The AN sending the Multicast Admission Control message MUST set the Result field to Ignore (0x0).

The AN MUST populate the ANCP Transaction Identifier field with a unique value, as described in Section 4.4.1 of [I-D.ietf-ancp-protocol] .

The AN MUST encode the Command TLV as specified in Section 4.3 with the following additional rules:

- the Accounting field MUST be set to 0;
- the Command Code field MUST be set to "0x01 - Add" when the message conveys a Join , to "0x02 - Delete" when the message conveys a Leave and to "0x03 - Delete All" when the message conveys a Leave of all channels (on the target);
- The Multicast-Flow TLV within the Command TLV identifies the multicast flow subject to the request for admission or release. When the Command Code is 0x03, the Multicast-Flow TLV is meaningless and MUST be omitted.
- the Request-Source-IP embedded TLV MAY be included by the AN to convey the IP address of the sender of the join/leave message (e.g. IGMP/MLD Join/Leave) that triggered the AN to include the

corresponding Command TLV in the Admission Control message. If it appears more than once, only the first instance is considered meaningful and the other instances are ignored.

- the Request-Source-MAC embedded TLV MAY be included by the AN to convey the MAC address of the sender of the join/leave message (e.g. IGMP/MLD Join/Leave) that triggered the AN to include the corresponding Command TLV in the Admission Control message. If it appears more than once, only the first instance is considered meaningful and the other instances are ignored.

4.4.2. Receiver Behaviour

On receipt of an Multicast Admission Control message, the NAS:

- MUST ignore the Result field;

- if the directive in the Multicast Admission Control message is "0x02 - Delete" or "0x03 - Delete All" and is processed correctly by the NAS, the NAS MUST NOT generate any ANCP message in response to the Multicast Admission Control message;

- if the directive in the Multicast Admission Control message is "0x01 - Add" and is accepted by the NAS, the NAS MUST generate a Multicast Replication Control in response to the Multicast Admission Control message. The Multicast Replication Control message:

  * MUST contain a Result set to Nack (0x1);

  * MUST contain a Transaction ID generated by the NAS (distinct non-zero, and linearly incremented by NAS for each request per adjacency);

  * MUST contain the directive as accepted by the NAS. The NAS MAY modify the Accounting field if flow accounting is required.

- if the directive in the Multicast Admission Control message is "0x01 - Add", is processed correctly but not accepted by the NAS (i.e. it does not pass the admission control or conditional access check), the NAS MAY generate a Multicast Replication Control message in response to the Multicast Admission Control message. This optional message can be used by the AN to maintain statistics about admission control rejections.

  In the future, the AN may be able to notify the subscriber that the request was rejected (e.g. using [I-D.morin-mboned-igmpmld-error-feedback]).

When used in this situation, the Multicast Replication Control message:

* MUST contain a Result set to 0x0;

* MUST contain a Transaction ID generated by the NAS (distinct non-zero, and linearly incremented by NAS for each request per adjacency);

* MUST contain the directive rejected by the NAS (i.e. Target TLV and Command TLV) but with a Command Code set to "0x04 - Admission Control Reject", "0x05 - Conditional Access Reject", or "0x06 - Admission Control and Conditional Access Reject".

- o if the Multicast Admission Control message cannot be processed correctly by the NAS (e.g. the message is malformed, the multicast flow does not exist etc.), the NAS MUST generate a Generic Response message (defined in Section 6.1.3 of [I-D.ietf-ancp-protocol]) with appropriate content indicating the reason for the failure.

4.5. Bandwidth Reallocation Request Message

The Bandwidth Reallocation Request message is used when the bandwidth delegation capability is included in the negotiated set. It MAY be sent either by the NAS or by the AN to request an adjustment in the amount of delegated bandwidth. It will be sent by the NAS typically to reduce the multicast bandwidth allocated to the AN in order for the NAS to satisfy a request to add one or more flows. Conversely, the AN will send a Bandwidth Reallocation Request to obtain additional bandwidth to satisfy a request to add a multicast channel. In each case, the requestor has a minimum requirement for additional bandwidth, and MAY ask for additional bandwidth beyond this amount (e.g., to handle anticipated future requests).

The Bandwidth Reallocation Request message contains two TLVs:

- o the Target TLV (Section 6.2.1 of [I-D.ietf-ancp-protocol] or an extension), specifying a single access line;

- o the Bandwidth-Request TLV (Section 5.8), specifying the required and preferred amounts of delegated bandwidth.

The Message Type for the Bandwidth Reallocation Request message is 146.

4.5.1. Sender Behaviour

The Result field in the header of the Bandwidth Reallocation Request message is not used and the sender MUST set it to Ignore (0x0).

The bandwidth values in the Bandwidth-Request TLV are expressed in terms of total multicast bandwidth allocated to the AN.

> The choice of "total bandwidth" rather than "incremental bandwidth" was made so that it would be easier for the AN and NAS to keep their respective views of the current amount of delegated bandwidth synchronized.

Because the values are totals rather than desired increments/ decrements, the relationship between the required amount and the preferred amount will differ depending on whether the Bandwidth Reallocation Request message is issued by the NAS or the AN.

- If the NAS is making the request, the preferred amount MUST be less than or equal to the required amount. The required amount MUST be less than the currently amount of delegated bandwidth.

- If the AN is making the request, the preferred amount MUST be greater than or equal to the required amount. The required amount MUST be greater than the currently amount of delegated bandwidth.

4.5.2. Receiver Behaviour

When the peer receives a valid Bandwidth Reallocation Request message, it SHOULD determine whether it can satisfy the request from its existing allocation of unused video bandwidth. If it decides that it can reallocate bandwidth to the peer, it MAY choose to return any amount between the required and the preferred amounts indicated in the Bandwidth Reallocation Request message.

The peer MUST return a Bandwidth Transfer message Section 4.6 indicating its decision. If the request is met, the Result field of the Bandwidth Transfer message MUST be set to Success (0x3), the Code field MUST be set to 0x000, and the Bandwidth-Allocation TLV (Section 5.5) MUST contain the new value of total multicast bandwidth. This new value MUST lie between the required and preferred values, inclusive, from the request message. If the request is not met, the Result field of the Bandwidth Transfer message MUST be set to Failure (0x4), the Code field MUST be set to 0x000, and the Bandwidth Allocation TLV MUST contain the value of the currently allocated amount of delegated bandwidth as the responder views it.

The following cases indicate that the sender holds a different view of the amount of delegated bandwidth from the receiver:

- the NAS receives a request where the required amount is less than its view of the current amount of delegated bandwidth;
- the AN receives a request where the required amount is greater than its view of the current amount of delegated bandwidth.

If one of these cases occurs, the receiver with one exception MUST send a Bandwidth Transfer message indicating Success.

- If the NAS received the request, the allocated amount in the NAS's response MUST be at least equal to NAS's view of the current amount of delegated bandwidth.
- If the AN received the request, the allocated amount in the AN's response MUST be no greater than the AN's view of the current amount of delegated bandwidth.

The exception is when the NAS receives a request while it has a request of its own outstanding. Handling of that case is described below.

While the cases just described are an error condition, the success response achieves a graceful recovery.

To avoid deadlock due to race conditions, the following rules MUST be applied:

a. If the NAS receives a Bandwidth Reallocation Request message while it has a Bandwidth Reallocation Request message of its own outstanding for the same access line, the NAS MUST provide an immediate failure response to the request from the AN, with a Code value set to 105 "Bandwidth request conflict".

b. If the AN receives a Bandwidth Reallocation Request message while it has a Bandwidth Reallocation Request message of its own outstanding for the same access line, the AN MUST release any bandwidth it has already committed to an outstanding Join request while it is awaiting a response from the NAS. It MUST decide upon and send its response to the NAS taking the released bandwidth into account.

If the receiver is unable to process the Bandwidth Reallocation Request message due to an error, then the receiver MUST return a Bandwidth Transfer message where:

- o the Result field is set to Failure (0x4),
- o the Code field is set appropriately to indicate the type of error that was detected,
- o the Bandwidth Allocation TLV contains the value of the current amount of delegated bandwidth as the responder views it, and
- o a Status-Info TLV MAY follow the Bandwidth Allocation TLV giving further information about the error.

This specification provides three new Code values applicable specifically to the contents of the Bandwidth-Request TLV. These Code values by their nature MUST only be used when the error is being reported in a Bandwidth Transfer message rather than a Generic Response message.

103 invalid preferred bandwidth amount. This indicates that the preferred and required amounts of bandwidth in the TLV do not have the numerical relationship described in the previous section.

104 inconsistent views of delegated bandwidth amount. This will appear only in a Bandwidth Transfer message from the NAS to the AN in the case where the NAS has an outstanding Bandwidth Reallocation Request. The recommended procedure for recovery is described in Section 4.6.2.

105 bandwidth request conflict. The NAS has rejected the AN's request for more bandwidth because the NAS has an outstanding bandwidth request.

4.6. Bandwidth Transfer Message

The Bandwidth Transfer message is used to transfer video bandwidth from the sender to the peer for a specific access line. This message MAY be sent either from the AN or from the NAS. As described in the previous section, it is the required response to a valid Bandwidth Reallocation Request message.

The Bandwidth Transfer message MAY also be used to transfer bandwidth autonomously from one peer to another. One example of this usage is to release bandwidth borrowed earlier by means of the Bandwidth Reallocation Request message. When the message is used in this way, the Result field in the Bandwidth Transfer message MUST be set to Ignore (0x0).

This allows the receiver to distinguish between an autonomous transfer and a response to a previous Bandwidth Reallocation

Request, for purposes of validation.

The Message Type for the Bandwidth Transfer message is 147. The Bandwidth Transfer message contains the following TLVs:

- the Target TLV, designating the access line concerned;
- an instance of the Bandwidth-Allocation TLV (Section 5.5). The bandwidth value in the Bandwidth-Allocation TLV is the new amount of delegated bandwidth allocated to the target.

4.6.1. Sender Behaviour

When sending a Bandwidth Transfer message where the Result value is Ignore (0x0) or Success (0x3), the following relationships MUST hold:

- if the message is sent by the NAS, the bandwidth value in the Bandwidth-Allocation TLV MUST be greater than or equal to the sender's view of the current amount of delegated bandwidth for the access line concerned;
- if the message is sent by the AN, the bandwidth value in the Bandwidth-Allocation TLV MUST be less than or equal to the sender's view of the current amount of delegated bandwidth for the access line concerned.

Further sender behaviour is specified above, in Section 4.5.2.

4.6.2. Receiver Behaviour

4.6.2.1. Behaviour of the NAS

If the amount of delegated bandwidth provided in the Bandwidth-Allocation TLV is not greater than the NAS's view of the current amount of delegated bandwidth, the NAS MUST update its view of the current amount of delegated bandwidth to the amount indicated in the Bandwidth Transfer message. This is required regardless of whether the Result field of that message indicates Success or Failure.

If the amount of delegated bandwidth provided in the Bandwidth-Allocation TLV is greater than the NAS's view of the current amount of delegated bandwidth, the NAS MAY accept the given value as its new value of delegated bandwidth. Alternatively, the NAS MAY force the AN to modify its view of the amount of delegated bandwidth to that held by the NAS, by sending a Port Management message for the target access line concerned, containing a Bandwidth-Allocation TLV with a value equal to the amount of delegated bandwidth the NAS wishes to enforce.

4.6.2.2. Behaviour of the AN

If the amount of delegated bandwidth provided in the Bandwidth-Allocation TLV of the Bandwidth Transfer message differs from the AN's view of the current amount of delegated bandwidth, the AN MUST update its view of the current amount of delegated bandwidth to the amount indicated in the Bandwidth Transfer message. This is required with the exception of a Bandwidth Transfer message with a Result field equal to Failure (0x4) and a Code field equal to 104 "Inconsistent views of delegated bandwidth amount" or 105 "Bandwidth request conflict". If Code value 104 is received, the AN MUST issue a Delegated Bandwidth Query Request message to determine the NAS's current view of the amount of delegated bandwidth. The AN MUST update its own view based on the value returned in the Delegated Bandwidth Query Response. If Code value 105 is received, the AN SHOULD carry out this procedure unless it can account for the discrepancy as a result of a transfer of bandwidth to the NAS that was carried out just before the incoming Bandwidth Transfer message was processed.

> The two Code values indicate a race condition where the AN may have just completed a transfer of bandwidth to the NAS. As a result, the value given in the Bandwidth Transfer message may be outdated, and the AN needs to query the NAS to find its latest view. The procedure assumes that ordering is preserved between the Bandwidth Transfer message sent by the AN in response to the NAS's request and the subsequent Delegated Bandwidth Query Request message.

If as the result of the procedures just described the AN determines that it has over-committed multicast bandwidth, it MUST NOT terminate any currently-active programs, but MUST NOT honour any more "join" requests until it is possible to do so within the limit set by its current value of delegated bandwidth.

4.7. Delegated Bandwidth Query Request Message

The Message Type for the Delegated Bandwidth Query Request (and Response) messages is 148.

The Delegated Bandwidth Query Request message MAY be sent either by the NAS or by the AN to retrieve the peer's view of the amount of delegated bandwidth. The request contains one TLV:

- a Target TLV designating the access line for which the information is requested.

4.7.1. Sender Behaviour

The sender MUST set the Result field in the header of the Delegated Bandwidth Query Request message to AckAll (0x2). The Code value MUST be set to 0x000. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in Section 4.4.1 of [I-D.ietf-ancp-protocol].

4.7.2. Receiver Behaviour

If the AN or NAS receives a valid Delegated Bandwidth Query Request message, it MUST respond with a Delegated Bandwidth Query Response message. The Result field in the header of the response MUST be set to Success (0x3). The Code field MUST be set to 0x000. The Transaction-Id field MUST be copied from the request message. The body of the response MUST contain the Target TLV, copied from the request message. Finally, the body of the response MUST contain a Bandwidth-Allocation TLV, containing the current amount of delegated bandwidth from the point of view of the receiver of the request.

If the contents of the Delegated Bandwidth Query Request message are in error, the receiver MUST return a Delegated Bandwidth Query Response message with the Result field in the header set to Failure (0x3). The Code field MUST be set to the value that indicates the nature of the error (e.g., 4 "Unrecognized target"). The Transaction-Id field MUST be copied from the request. The body of the response MUST contain the Target TLV copied from the request. This MAY be followed by a Status-Info TLV giving further information about the error.

4.8. Delegated Bandwidth Query Response Message

The Delegated Bandwidth Query Response message is sent in reply to a Delegated Bandwidth Query Request. The response to a valid request contains two TLVs:

- the Target TLV, copied from the request;
- a Bandwidth-Allocation TLV, giving the responder's view of the current amount of multicast bandwidth delegated to the AN.

The Message Type for the Delegated Bandwidth Query Response message is 148.

4.8.1. Sender Behaviour

Sender behaviour for the Delegated Bandwidth Query Response message is specified in Section 4.7.2.

4.8.2. Receiver Behaviour

If the Delegated Bandwidth Query Response message indicates Success (0x3), the following actions apply.

4.8.2.1. Behaviour at the NAS

If the amount of delegated bandwidth provided in the Bandwidth-Allocation TLV is less than the NAS's view of the current amount of delegated bandwidth, the NAS MUST update its view of the current amount of delegated bandwidth to the amount indicated in the Delegated Bandwidth Query Response message.

If the amount of delegated bandwidth provided in the Bandwidth-Allocation TLV is greater than the NAS's view of the current amount of delegated bandwidth, the NAS MAY accept the given value as its new value of delegated bandwidth. Alternatively, the NAS MAY force the AN to modify its view of the amount of delegated bandwidth to that held by the NAS, by sending a Port Management message for the target access line concerned, containing a Bandwidth-Allocation TLV with a value equal to the amount of delegated bandwidth the NAS wishes to enforce.

4.8.2.2. Behaviour at the AN

The AN SHOULD accept the value returned in the Bandwidth-Allocation TLV of the Delegated Bandwidth Query Response message as the correct value of the current amount of delegated bandwidth. If the AN has currently committed more than this amount to active programs, it MUST NOT cease replicating the flows concerned, but MUST NOT honour any more Join requests until possible to do so within the new limit.

> A race condition is possible, where the AN sends a query, the NAS requests more bandwidth, then receives and responds to the query, then receives the Bandwidth Transfer message responding to its request. It is up to the AN to take appropriate action in this case. The best action appears to be not to act on the result of the first query, but to repeat the query after sending the Bandwidth Transfer message. Similar considerations apply to a race between queries from both sides.

4.9. Multicast Flow Query Request and Response Messages

This section defines two new messages called the Multicast Flow Query Request and Multicast Flow Query Response. The Multicast Flow Query Request is sent by the NAS to request information about the multicast flows that are active on the AN. The Multicast Flow Query Response is sent in response by the AN to provide the requested information to

the NAS.

The Message Type for the Multicast Flow Query Request and Multicast Flow Query Response messages is 149.

The contents of the Multicast Flow Query Request and Response depend on the nature of the query, as described below.

4.9.1. Sender Behaviour

The sender of a Multicast Flow Query Request message MUST set the Result field to AckAll (0x2). The Code field MUST be set to 0x000. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in section 4.4.1 of [I-D.ietf-ancp-protocol].

The Multicast Flow Query Request MAY be used by the NAS to retrieve:

- the AN's view of which multicast flows are currently active on a specified set of access ports; or
- the AN's view of the access ports on which a specified set of multicast flows are currently active; or
- the AN's view of all the multicast flows currently active on each and every port of the AN.

To retrieve the AN's view of which multicast flows are currently active on a given port of the AN, the NAS MUST include a Target TLV in the Multicast Flow Query Request payload identifying that port. The Target TLV is encoded as specified in [I-D.ietf-ancp-protocol].

To retrieve the AN's view of the ports currently receiving a given multicast flow, the NAS MUST include a Multicast-Flow TLV in the Multicast Flow Query Request payload identifying that flow. The Multicast-Flow TLV is encoded as specified in Section 5.11.

The NAS MAY include multiple Target TLVs or multiple Multicast-Flow TLVs in the Multicast Flow Query Request, but MUST NOT include both Target and Multicast-Flow TLVs in the same message.

To retrieve the AN's view of all of the multicast flows currently active on each port of the AN, the NAS MUST send a Multicast Flow Query Request which does not contain any instance of the Target TLV or the Multicast-Flow TLV.

4.9.2. Receiver Behaviour

The AN MUST respond to a Multicast Flow Query Request message that has a valid format and a valid content with a Multicast Flow Query Response message. The Result field in the response MUST be set to Success (0x3). The Code field MUST be set to 0x000. The Transaction-Id field MUST be copied from the request.

If the Multicast Flow Query Request contained one (or more) Target TLVs, the AN MUST include, for each of these Target TLVs, the following set of TLVs:

- Target TLV. This MUST be identical to the Target TLV in the received Multicast Flow Query Request message.

- Multicast-Flow TLV(s). The Multicast-Flow TLV MUST appear once per multicast flow that is currently active on the AN port identified in the preceding Target TLV.

The Target TLVs MUST appear in the response from the AN in the same order as in the query from the NAS.

If the Multicast Flow Query Request contained one (or more) Multicast-Flow TLVs, the AN MUST include, for each of these Multicast-Flow TLVs, the following set of TLVs:

- Multicast-Flow TLV. This MUST be identical to the Multicast- Flow TLV in the received Multicast Flow Query Request message.

- Target TLV(s). The Target TLV MUST appear once per AN port on which the multicast flow identified in the preceding Multicast Flow TLV is active.

The Multicast-Flow TLVs MUST appear in the response from the AN in the same order as in the query from the NAS.

If the Multicast Flow Query Request contained no Target TLV and no Multicast Flow TLV, the AN MUST include, for each AN port currently receiving multicast flow(s), the following set of TLVs:

- Target TLV. This MUST identify one AN port.

- Multicast-Flow TLV(s). The Multicast-Flow TLV MUST appear once per Multicast Flow that is currently active on the AN port identified in the preceding Target TLV.

If the contents of the Multicast Flow Query Request are in error, the AN MUST reply with a Multicast Flow Query Response message with the

Result field set to Failure (0x4) and the Code field set to indicate the nature of the error. If the request contained multiple instances of the Target TLV or the Multicast-Flow TLV and one of these is in error, the response message MUST contain the results for the preceding instances of the TLV as if there had been no error. These successful results MUST be followed by the TLV in error, copied from the request. The AN MUST NOT do further processing of the request. The AN MAY add a Status-Info TLV to provide further information on the nature of the error.

4.10. Committed Bandwidth Report Message

This section describes the Committed Bandwidth Report message, which is sent from the AN to the NAS to report the most recent amount of multicast bandwidth usage committed to one or more access lines.

The Message Type for the Committed Bandwidth Report message is 150.

The Committed Bandwidth Report message contains one or more instances of the Committed-Bandwidth TLV, as described in Section 5.13.

4.10.1. Sender Behaviour

The sender of a Committed Bandwidth Report message MUST set the Result field to Ignore (0x0). The Code field MUST be set to 0x000. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in section 4.4.1 of [I-D.ietf-ancp-protocol].

Each instance of the Committed-Bandwidth TLV included in the message MUST identify an access line for which the amount of committed multicast bandwidth has changed since the previous Committed Bandwidth Report message was sent and MUST report the latest amount of multicast bandwidth committed to that line. There MUST be only one instance of the Committed-Bandwidth TLV present in the message for any given access line. The message MUST include an instance of the Committed-Bandwidth TLV for every access line for which committed multicast bandwidth has changed since the previous Committed Bandwidth Report message was sent.

Further behaviour at the AN is specified in Section 6.2.2.

4.10.2. Receiver Behaviour

The usage of the contents of a Committed Bandwidth Report message received by the NAS is implementation-dependent. One example is that the NAS uses the reports of multicast bandwidth commitments to adjust its forwarding scheduler operation to provide the intended level of

QoS.

The NAS MUST NOT reply to a valid Committed Bandwidth Report message. The NAS MAY send a Generic Response message indicating the nature of any errors detected in a Committed Bandwidth Report message that it has received.

5. ANCP TLVs For Multicast

This section defines new ANCP TLVs for the control of multicast flows.

5.1. Multicast-Service-Profile TLV

This document defines the new Multicast-Service-Profile TLV.

The Multicast-Service-Profile TLV MAY be included in a Provisioning message as specified in Section 4.1.

The Multicast-Service-Profile TLV is illustrated in Figure 6. It consists of a TLV header encapsulating a single instance of the Multicast-Service- Profile-Name TLV and one or more instances of the List-Action TLV.

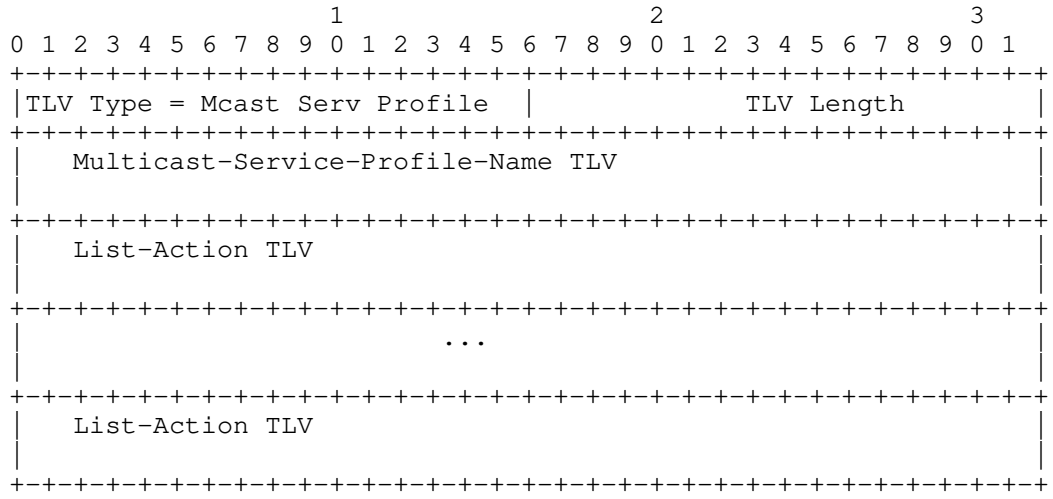

Figure 6: Multicast-Service-Profile TLV

The Multicast-Service-Profile TLV has the following fields:

- The Multicast-Service-Profile TLV Type is 0x13.
- The TLV length is determined by the contents following the TLV header.
- The Multicast-Service-Profile-Name TLV is described in Section 5.2. The Multicast-Service-Profile-Name TLV MUST contain an identifier which is unique over all profiles provisioned to the same AN partition. This identifier will be used to refer to the

profile when activating it for a given target within a Port Management message (see Section 4.2).

- The List-Action TLV is described in Section 5.3. The List-Action TLV(s) provide the content of a newly defined multicast service profile or modify the existing content. If more than one List-Action TLV is present, the order of the TLVs may be significant, since List- Action TLVs are processed in the order in whch they appear.

5.2. Multicast-Service-Profile-Name TLV

The Multicast-Service-Profile-Name TLV carries the identifier of a multicast service profile provisioned on the AN. It is illustrated in Figure 7.

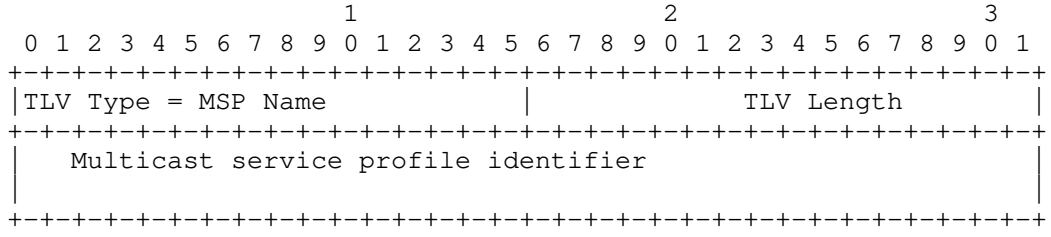

Figure 7: Multicast-Service-Profile-Name TLV

The Multicast-Service-Profile-Name TLV has the following fields:

- The Multicast-Service-Profile-Name TLV Type is 0x18.
- TLV Length: up to 255 octets.
- Multicast service profile identifier: an opaque sequence of bits identifying a specific multicast service profile.

  The identifier could have the form of human-readable text or an arbitrary binary value, depending on the operator's practices.

5.3. List-Action TLV

The List-Action TLV identifies multicast flows to be added to or removed from a list of White-, Black-, or Grey-listed flows. It is meaningful only in association with a Multicast-Service-Profile-Name TLV identifying the profile to which the List-Action TLV applies. Such an association can be achieved by placing both TLVs in the same base message payload or as embedded TLVs of another TLV such as the Multicast-Service-Profile. The List-Action TLV is shown in Figure 8.

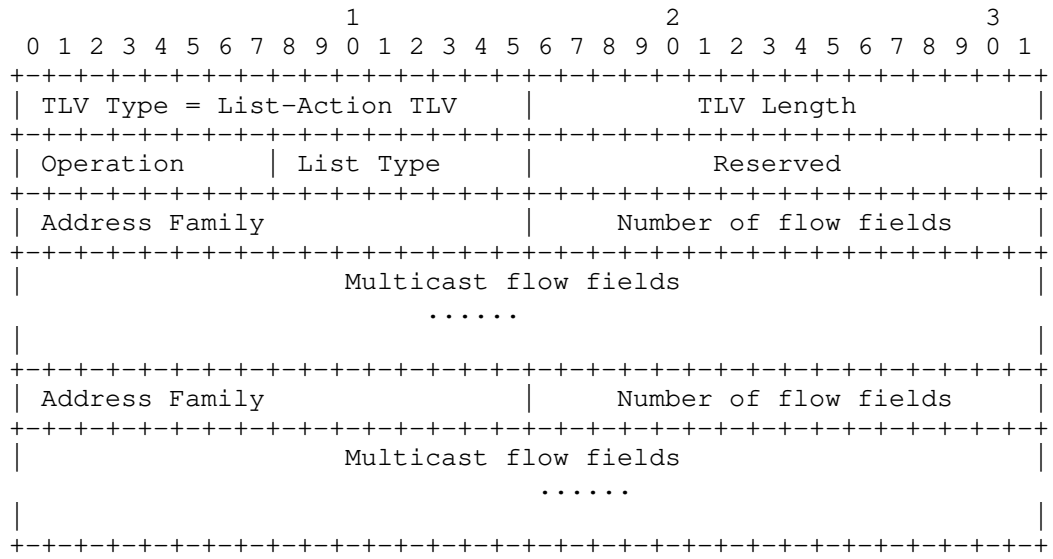

Figure 8: List-Action TLV

The List-Action TLV contains the following fields:

- The List-Action TLV Type is 0x21.
- TLV Length: length of the subsequent contents.
- Operation: operation to be performed upon the White, Black, or Grey list identified by the List Type field within the profile identified by the associated Multicast-Service-Profile-Name embedded TLV. The possible values are:
  - Add (0x01): the multicast flow fields are to be added to the list.
  - Delete (0x02): the multicast flow fields are to be removed from the list. Each multicast flow field in the List-Action MUST match exactly an existing entry in the list concerned. Thus to remove part of the range provided by a wildcarded list entry, it is necessary to remove the entire entry and add back the remaining partial range(s).
  - Replace (0x03): the multicast flow fields replace the existing contents of the list.

- List Type: the list type being modified by this List-Action. The possible values are White (0x01), Black (0x02), or Grey (0x03).

- Reserved: a sender MUST set this field to 0x0000. A receiver MUST ignore the contents of this field.

- Address Family: the IP version of the set of multicast flow fields that follow, encoded according to [PIMreg]. Possible values are 0x0001 (IPv4) or 0x0002 (IPv6). Either an IPv4 list or an IPv6 list or both MAY be present in the List-Action TLV.

- Number of flow fields: the number of multicast flow fields of the given address family which follow.

- Multicast flow field: a field identifying one or more multicast flows. It consists of an 8-bit group address prefix length, an 8-bit source address prefix length, a 0-16 octet group prefix, and a 0-16 octet source prefix, as shown in Figure 9.

Each multicast flow field refers either to a Source-Specific Multicast (SSM) channel or to an Any Source Multicast (ASM) group. The scope of the designation may be broadened to multiple channels or groups through use of prefix length values smaller than the total address length for the given address family. Multicast flow fields MUST be placed consecutively within the embedded TLV without intervening padding except to round out individual addresses to the nearest octet boundary.

A multicast flow field consists of two single-octet prefix lengths followed by zero to two prefix values as shown in Figure 9:

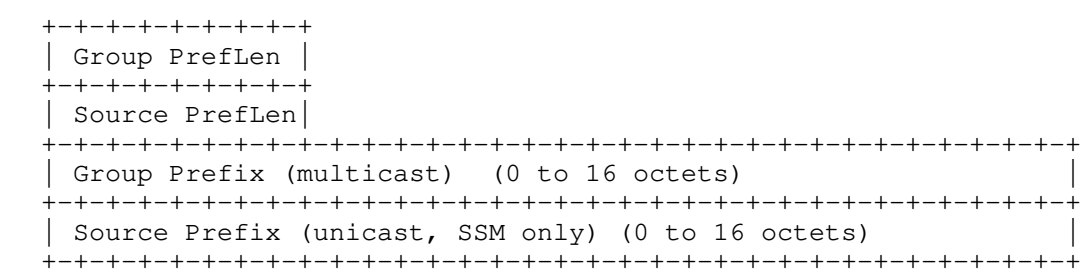

Figure 9: Organization of a Single Multicast Flow Field

The prefix length has its usual meaning. It is the number of most-significant bits specified within the corresponding prefix. The prefix length MAY vary from 0 to 32 in the IPv4 sub-list, and from 0 to 128 in the IPv6 sub-list.

A value of 0x00 for either the Group PrefLen (prefix length) or the Source PrefLen indicates that any value of the corresponding address will match (wild card). If the value 0x00 is provided for a particular prefix length, the corresponding prefix MUST be omitted from the field contents. In particular, a value of 0x00 for the Source PrefLen indicates an ASM multicast entry, and the Source Prefix will be absent.

The length of a Source or Group Prefix field is equal to (PrefLen + 7)/8 octets, truncated to the nearest integer. Unused bits at the end of the prefix MUST be set to zeroes.

5.4. Sequence-Number TLV

The Sequence-Number TLV conveys a sequence number of some sort. The specific meaning of the sequence number is message-specific. Within this specification, the Sequence-Number TLV is used as a embedded TLV within a Status-Info TLV, in a Generic Response reporting a failed command within a Multicast Replication Control or Multicast Admission Request message. It identifies the sequence number within the message of the command that failed.

The Sequence-Number TLV has the format shown in Figure 10.

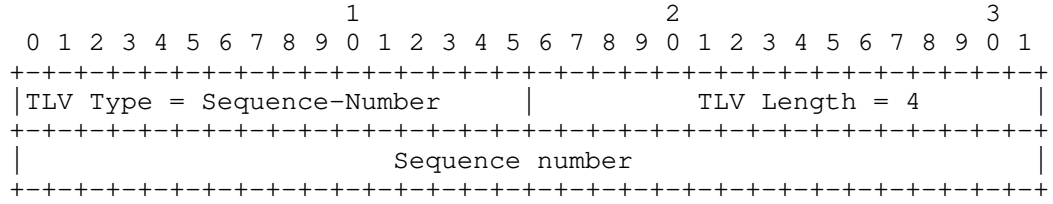

Figure 10: Sequence-Number TLV

The Sequence-Number TLV has the following fields:

- The Sequence-Number TLV Type is 0x22.
- TLV length is 0x0004.
- Sequence number: the sequence number of a specific entity within a series, where numbering starts from 1 for the first entity in the series. Represented as a 32-bit binary number, most significant bit first.

5.5. Bandwidth-Allocation TLV

The Bandwidth-Allocation TLV is used to indicate the total amount of video bandwidth delegated to the AN for multicast admission control for a given access line, in kilobits per second. The TLV has the format shown in Figure 11.

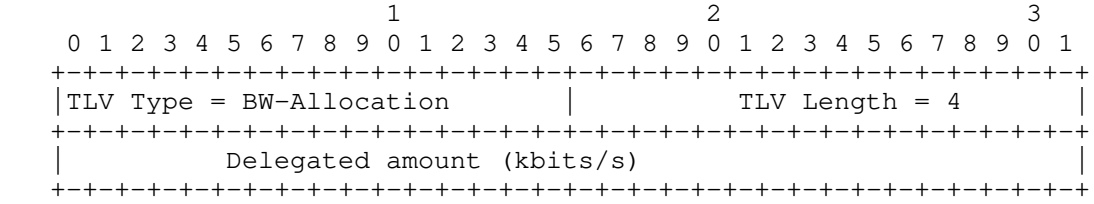

Figure 11: The Bandwidth-Allocation TLV

The Bandwidth-Allocation TLV has the following fields:

- The Bandwidth-Allocation TLV Type is 0x15.
- TLV length is 4.
- Delegated amount: the bandwidth amount delegated to the AN for admission of multicast video on a given port, kilobits per second. Presented as a 32-bit binary value, most significant bit first.

5.6. White-List-CAC TLV

The White-List-CAC TLV is used to indicate that the NAS wishes the AN to do admission control for White-listed flows. Details on when the White-List-CAC TLV may be provisioned are specified in Section 6. The White- List-CAC TLV is illustrated in Figure 12.

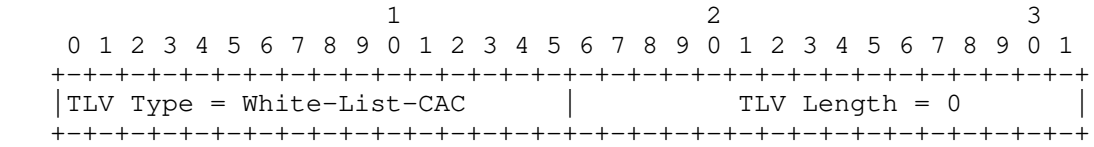

Figure 12: White-List-CAC TLV

The White-List-CAC TLV contains the following fields:

- The White-List-CAC TLV Type is 0x24.
- TLV length is 0, since the TLV contains no data other than the TLV header.

5.7. MRepCtl-CAC TLV

The MRepCtl-CAC TLV is used to indicate that the NAS wishes the AN to do admission control for flows added by the Multicast Replication Control message. Details on when the MRepCtl-CAC TLV may be provisioned are specified in Section 6. The MRepCtl-CAC TLV is illustrated in Figure 13.

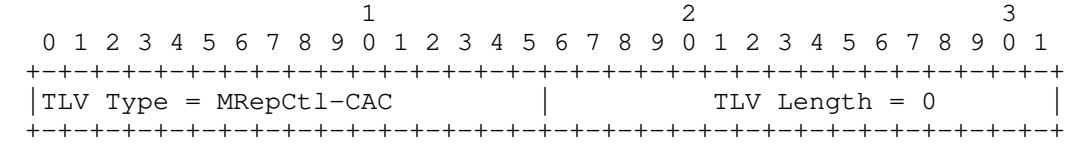

Figure 13: MRepCtl-CAC TLV

The MRepCtl-CAC TLV contains the following fields:

- The MRepCtl-CAC TLV Type is 0x25.
- TLV length is 0, since the TLV contains no data other than the TLV header.

5.8. Bandwidth-Request TLV

The Bandwidth-Request TLV is used to request an adjustment of the total amount of video bandwidth allocated to the AN for multicast admission control for a given line. The "Required amount" field indicates the minimum adjustment required to meet the request. The "Preferred amount" field indicates the adjustment the requestor would prefer to have, if possible. Section 4.5 discusses the required relationships between the "Required amount", "Preferred amount", and current values of total bandwidth allocated to the AN.

The Bandwidth-Request TLV has the format shown in Figure 14.

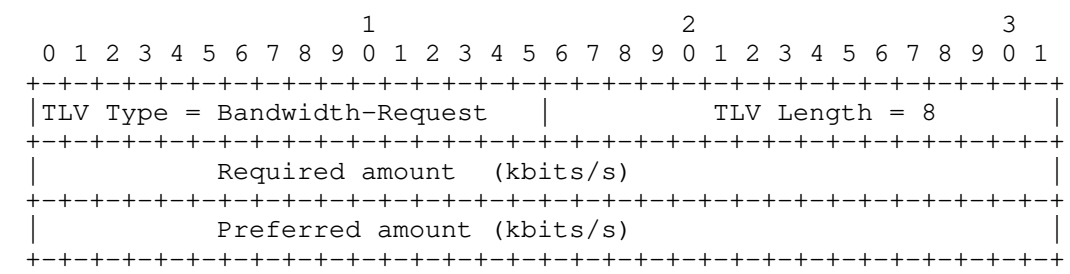

Figure 14: The Bandwidth-Request TLV

The Bandwidth-Request TLV has the following fields:

- The Bandwidth-Request TLV Type is 0x16.
- The TLV length is 8 octets.
- Required amount: the minimum or maximum amount, depending on whether the sender is the AN or the NAS respectively, of delegated video bandwidth that is being requested, in kilobits per second. Presented as a 32-bit binary value, most significant bit first.
- Preferred amount: the preferred amount of delegated video bandwidth that is being requested, in kilobits per second. Presented as a 32-bit binary value, most significant bit first.

5.9. Request-Source-IP TLV

The Request-Source-IP TLV provides the IP address of the entity that originated a specific request to join or leave a multicast channel. The TLV is illustrated in Figure 15.

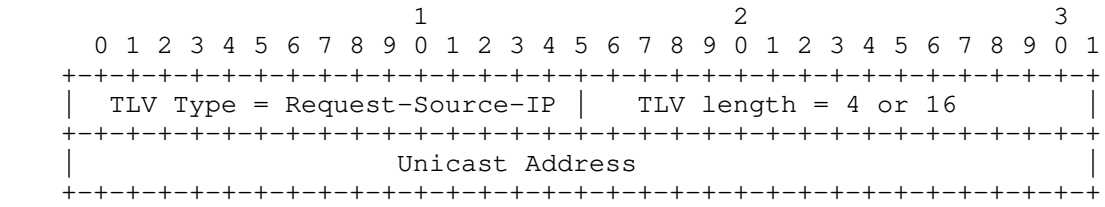

Figure 15: Request-Source-IP TLV

The Request-Source-IP TLV contains the following fields:

- The Request-Source-IP TLV Type is 0x92.
- TLV length is 4 for an IPv4 address or 16 for an IPv6 address.
- Unicast address: IP address of the source of a multicast flow join request, in network byte order.

5.10. Request-Source-MAC TLV

The Request-Source-MAC TLV provides the MAC address of the entity that originated a specific request to join or leave a multicast channel. The TLV is illustrated in Figure 16.

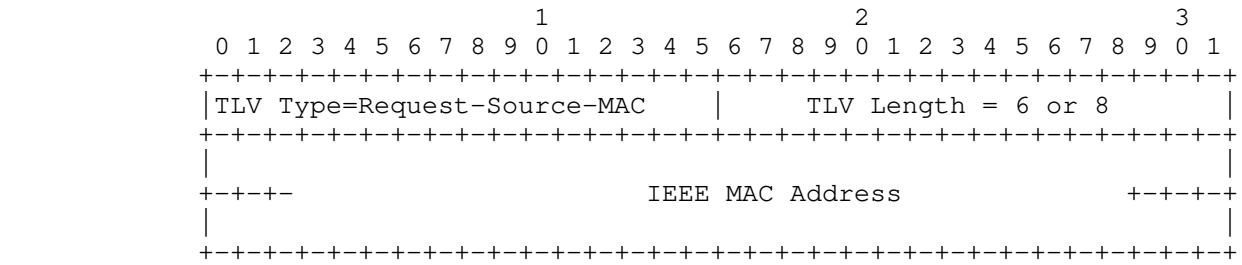

Figure 16: Request-Source-MAC TLV

The Request-Source-MAC TLV contains the following fields:

- The Request-Source-MAC TLV Type is 0x93.
- TLV length is either 6 octets (MAC-48 or EUI-48) or 8 octets (EUI-64).
- IEEE MAC Address: MAC address of the device originating the request to join a multicast flow. Within the address, bytes and bits respectively shall be ordered from most to least significant, consistently with [IEEE48] for MAC-48 and EUI-48, and with [IEEE64] for EUI-64.

  EUI-48 and EUI-64 are registered trademarks of the IEEE.

5.11. Multicast-Flow TLV

The Multicast-Flow TLV specifies a multicast flow in terms of its multicast group address and, if applicable, its unicast source address. It is illustrated in Figure 17.

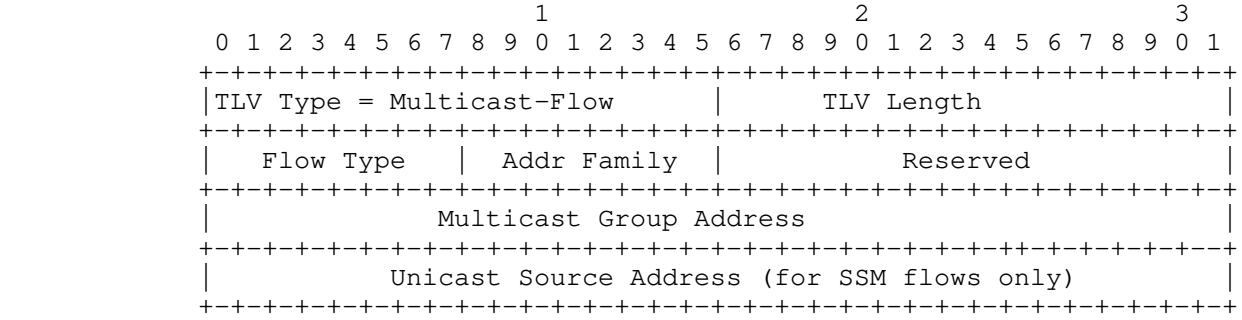

Figure 17: Multicast-Flow TLV

The Multicast-Flow TLV has the following fields:

- The Multicast-Flow TLV Type is 0x19.
- TLV Length: ranges from a minimum of 8 (for an ASM IPv4 flow) to 36 (for an IPv6 SSM flow.
- Flow Type: 0x01 for Any Source Multicast (ASM), 0x02 for Specific-Source Multicast (SSM).
- Addr Family: address family of the multicast source and group addresses, encoded in accordance with the IANA PIM Address Family registry ([PIMreg]). 0x01 indicates IPv4, 0x02 indicates IPv6.
- Reserved: MUST be set to 0x0000 by the sender and MUST be ignored by the receiver.

  One possible use for this field in the future is to indicate the presence of PIM Join attributes attached to the source address (see [RFC5384]). The applicability of PIM attributes in the context of ANCP is for further study.

- Multicast Group Address: a multicast group address within the given address family. The group address MUST always be present.
- Unicast Source Address: unicast address within the given address family. If the Flow Type is 0x01 (ASM), a source address MUST NOT be present. If the Flow Type is 0x02 (SSM), a source address MUST be present.

5.12. Report-Buffering-Time TLV

The Report-Buffering-Time TLV provides the time for which a Committed Bandwidth Report message must be held with the intention of accumulating multiple reports of changed committed multicast bandwidth in one report, to reduce the volume of messages sent to the NAS. For further information see Section 6.2.2. The TLV is illustrated in Figure 18.

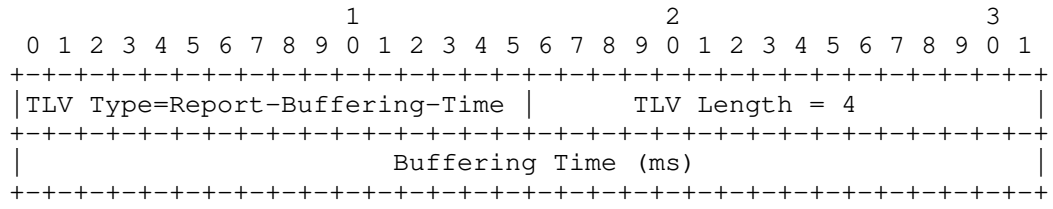

Figure 18: Report-Buffering-Time TLV

The Report-Buffering-Time TLV contains the following fields:

- The Report-Buffering-Time TLV Type is 0x94.
- TLV length is 4 octets.
- Buffering Time is a 32-bit unsigned integer containing a time value in ms.

5.13. Committed-Bandwidth TLV

The Committed-Bandwidth TLV identifies an access line and provides the current amount of multicast bandwidth that the AN has committed to it. The TLV is illustrated in Figure 19.

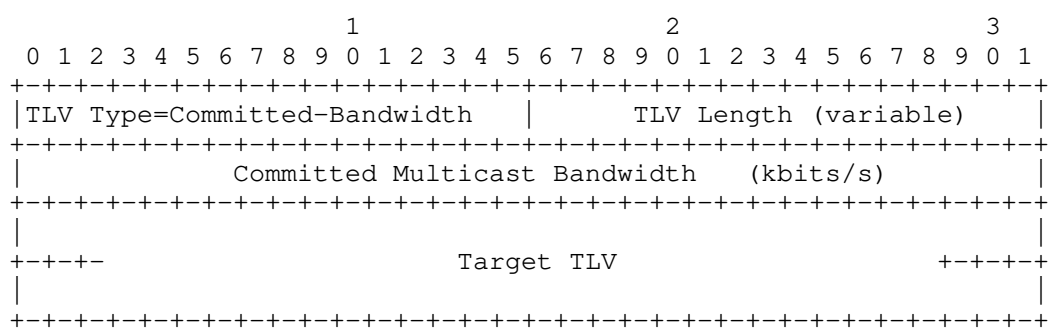

Figure 19: Committed-Bandwidth TLV

The Committed-Bandwidth TLV contains the following fields:

- The Committed-Bandwidth TLV Type is 0x95.
- TLV length is 4 octets plus the length of the Target TLV including its header and any padding.
- Committed Multicast Bandwidth is a 32-bit unsigned integer providing a bandwidth amount in kbits/s.
- The Target TLV identifies the access line to which this amount of multicast bandwidth is currently committed.

6. Multicast Capabilities

Section 4.3 of [I-D.ietf-ancp-protocol] defines a capability negotiation mechanism as well as a number of capabilities. This section defines five new capabilities in support of different modes of multicast operation:

- NAS-initiated replication (capability type 0x03);
- committed multicast bandwidth reporting (capability type 0x05);
- conditional access with white and black lists (capability type 0x06);
- conditional access with grey lists (capability type 0x07);
- bandwidth delegation (capability type 0x08).

The "Capability Data" field within the Capability TLV for all of these capabilities is empty. All of these capabilities are independent of the access technology.

The remainder of this section consists of three sub-sections. Section 6.1 specifies the protocol elements that must be implemented in order to support each capability. Section 6.2 specifies the procedures that apply to each capability on its own. Section 6.3 specifies how the capabilities interact if more than one multicast capability is included in the set of capabilities negotiated between the AN and the NAS.

Note that if a request contains content that is not supported (according to the tables in Section 6.1) by the negotiated set of multicast capabilities, the appropriate response is to return a Generic Response message indicating Failure (0x4) with an appropriate code value (e.g., 84 "TLV or value not supported by negotiated capability set"). The body of the message MUST contain a Status-Info TLV. See Sections 6.1.3 and 6.2.3 in [I-D.ietf-ancp-protocol] for more details.

6.1. Required Protocol Support

This section specifies the protocol elements that MUST be implemented to support each of the four multicast capabilities. Support of multiple multicast capabilities requires implementation of the union of the sets of protocol elements applying to each of the individual capabilities in the supported set.

6.1.1. Protocol Requirements For NAS-initiated Replication

Table 1 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the NAS-initiated replication multicast capability.

Reference	Protocol Element
Section 4.1	Provisioning message with MRepCtl-CAC TLV
Section 4.2	Port Management message with Bandwidth-Allocation TLV.
Section 4.3	Multicast Replication Control message
Section 4.9	Multicast Flow Query Request and Response messages
Section 5.4	Command Number TLV
Section 5.7	MRepCtl-CAC TLV
Section 5.11	Multicast-Flow TLV

Table 1: Protocol Requirements For NAS-initiated Replication

6.1.2. Protocol Requirements For Committed Multicast Bandwidth Reporting

Table 2 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the NAS-initiated replication multicast capability.

Reference	Protocol Element
Section 4.1	Provisioning message with Report-Buffering-Time TLV
Section 4.10	Committed Bandwidth Report message
Section 4.9	Multicast Flow Query Request and Response messages
Section 5.12	Report-Buffering-Timer TLV
Section 5.13	Committed-Bandwidth TLV
Section 5.11	Multicast-Flow TLV

Table 2: Protocol Requirements For Committed Multicast Bandwidth Reporting

6.1.3. Protocol Requirements For Conditional Access With White and Black Lists

Table 3 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the conditional access with white and black lists multicast capability.

Reference	Protocol Element
Section 4.1	Provisioning message with Multicast-Service-Profile TLV, White and Black lists only, and White-List-CAC TLV
Section 4.2	Port Management message with Multicast-Service-Profile-Name and Bandwidth-Allocation TLVs.
Section 4.9	Multicast Flow Query Request and Response messages
Section 5.1	Multicast-Service-Profile TLV
Section 5.2	Multicast-Service-Profile-Name TLV
Section 5.3	List-Action TLV, White and Black lists only
Section 5.5	Bandwidth-Allocation TLV
Section 5.6	White-List-CAC TLV
Section 5.11	Multicast-Flow TLV

Table 3: Protocol Requirements For Conditional Access with White and Black Lists

6.1.4. Protocol Requirements For Conditional Access With Grey Lists

Table 4 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the conditional access with grey lists multicast capability.

Reference	Protocol Element
Section 4.1	Provisioning message with Multicast-Service-Profile TLV, Grey lists only, and MRepCtl-CAC TLV
Section 4.2	Port Management message with Multicast-Service-Profile-Name and Bandwidth-Allocation TLVs.
Section 4.3	Multicast Replication Control message
Section 4.4	Multicast Admission Control Message
Section 4.9	Multicast Flow Query Request and Response messages
Section 5.1	Multicast-Service-Profile TLV, Grey lists only
Section 5.2	Multicast-Service-Profile-Name TLV
Section 5.3	List-Action TLV, Grey lists only
Section 5.4	Command Number TLV
Section 5.5	Bandwidth-Allocation TLV
Section 5.7	MRepCtl-CAC TLV
Section 5.9	Request-Source-IP TLV
Section 5.10	Request-Source-MAC TLV
Section 5.11	Multicast-Flow TLV

Table 4: Protocol Requirements For Conditional Access with Grey Lists

6.1.5. Protocol Requirements For Delegated Bandwidth

Table 5 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the delegated bandwidth multicast capability.

Reference	Protocol Element
Section 4.2	Port Management message with Bandwidth-Allocation TLV.
Section 4.5	Bandwidth Reallocation Request Message
Section 4.6	Bandwidth Transfer Message
Section 4.7	Delegated Bandwidth Query Request Message
Section 4.8	Delegated Bandwidth Query Response Message
Section 4.9	Multicast Flow Query Request and Response messages
Section 5.5	Bandwidth-Allocation TLV
Section 5.8	Bandwidth-Request TLV
Section 5.11	Multicast-Flow TLV

Table 5: Protocol Requirements For Delegated Bandwidth

6.2. Capability-Specific Procedures for Providing Multicast Service

This section describes multicast service procedures for each capability as if it were the only multicast capability within the negotiated set. Procedures involving combinations of multicast capabilities are described in Section 6.3.

The use of the Multicast Flow Query Request and Response messages to determine the association between multicast flows and ports is common to all multicast capabilities. No additional text is required here, beyond that already given in Section 4.9 to describe the use of those messages.

6.2.1. Procedures For NAS-Initiated Replication

NAS-initiated replication MAY be negotiated to support a mode of operation where IGMP/MLD requests are terminated on the NAS. Alternatively, it MAY be negotiated to allow the NAS to respond to requests sent by other means (e.g., through application signalling) that require the replication of multicast channels to a given access line.

6.2.1.1. Provisioning

The NAS MAY perform admission control for NAS-initiated replication. In this case, it MUST NOT include the MRepCtl-CAC TLV in a Provisioning message sent to the AN. Alternatively, the NAS MAY enable admission control at the AN for NAS-initiated replication. To do this, it MUST include the MRepCtl-CAC TLV in a Provisioning message sent to the AN and it MUST also include a Bandwidth-Allocation TLV in a Port Management message for each access line.

6.2.1.2. Multicast Service Procedures

The procedures associated with NAS-initiated replication are straightforward. To initiate replication, the NAS MUST send a Multicast Replication Control message to the AN, containing one or more commands adding flows, as described in Section 4.3.1. To terminate replication the NAS MUST send a Multicast Replication Control message where the commands delete instead of adding the flows. The AN acts upon these messages as specified in Section 4.3.2.

6.2.2. Procedures For Committed Bandwidth Reporting

Committed bandwidth reporting MAY be negotiated if the NAS requires current knowledge of the amount of multicast bandwidth committed to each access line and cannot obtain this information by other means.

6.2.2.1. Provisioning

The default buffering time when committed bandwidth reporting is enabled is zero (immediate reporting). To change this, the NAS MAY send an instance of the Report-Buffering-Time TLV containing a non-zero time value to the AN in a Provisioning message. If the NAS subsequently wishes to change the buffering time again, it MAY do so in another Provisioning message.

6.2.2.2. Multicast Service Procedures

If the buffering time for committed bandwidth reporting is zero, the AN MUST send a Committed Bandwidth Report message to the NAS each time the amount of multicast bandwidth committed to any access line under its control changes.

If a non-zero value is provided in the Report-Buffering-Time TLV, the AN at any given moment is in one of two states: not-buffering, or buffering. The AN enters buffering state if it is in not-buffering state and the multicast bandwidth amount committed to some access line changes. It leaves buffering state when the AN sends a

Committed Bandwidth Report.

Upon entry to the buffering state, the AN MUST start a buffering timer and create a Committed Bandwidth Report message containing a Committed-Bandwidth TLV for the triggering access line, but MUST NOT send it. If a multicast bandwidth change occurs for another access line, the AN MUST add a new Committed-Bandwidth TLV to the message for that additional line. If a multicast bandwidth change occurs for a line for which a Committed-Bandwidth TLV is already present in the buffered report, the AN MUST update the Committed-Bandwidth TLV to contain the new bandwidth value, rather than adding another Committed-Bandwidth TLV for the same access line.

The buffering timer expires after the period provided by the Report-Buffering-Time TLV. When it expires, the AN MUST send the Committed Bandwidth Report message that it has been accumulating to the NAS. Exceptionally, the AN MAY choose to send the message before the timer expires, in which case it MUST clear the buffering timer when the message is sent. In either case, the AN enters the not-buffering state as a result.

Report buffering implies that NAS reaction to changes in multicast bandwidth usage is delayed by the amount of the buffering period. The choice of buffering period must take this into consideration.

6.2.3. Procedures For Conditional Access With Black and White Lists

6.2.3.1. Provisioning

The NAS provisions named multicast service profiles containing White and Black lists on the AN using the Provisioning message containing one or more Multicast-Service-Profile TLVs. The NAS MAY update the contents of these profiles from time to time as required, by sending additional Provisioning messages with Multicast-Service-Profile TLVs containing incremental modifications to the existing White and Black lists or replacements for them.

The NAS assigns a specific multicast service profile to an individual access line using the Port Management message containing a Multicast-Service-Profile-Name TLV. The NAS MAY change the multicast service profile for a given access line at any time by sending a Port Management message identifying a new multicast service profile.

The NAS MAY choose to enable admission control at the AN for White-listed flows. To do this, it MUST send a Provisioning message as described in Section 4.1, which includes the White-List-CAC TLV and it MUST provide a multicast bandwidth allocation for each access line by including a Bandwidth-Allocation TLV in a Port Management message.

6.2.3.2. Multicast Service Procedures

The conditional access with White and Black lists capability assumes that IGMP/MLD requests are terminated on the AN. When the AN receives a "join" request, it MUST check to see whether the requested flow is White-listed or Black-listed as described below. Requests for Black-listed flows MUST be discarded. If the NAS has enabled admission control on the AN as described in the previous section, but a White-listed flow would cause the amount of committed multicast bandwidth to exceed the provisioned limit, the request MUST be discarded. The AN replicates flows passing these checks to the access line.

To determine if a requested flow is White-listed, the AN searches for a best match to the flow in the applicable multicast service profile. Matching is done on the prefixes specified in the profile, ignoring the address bits of lower order than those in the prefix.

If the requested multicast flow matches multiple lists associated with the access line, then the most specific match will be considered by the AN. If the most specific match occurs in multiple lists, the Black list entry takes precedence over the White list. In this context, the most specific match is defined as:

- first, most specific match (longest prefix length) on the multicast flow address (i.e., on G of <S,G>)

- then, most specific match (longest prefix length) on the unicast source address (i.e. on S of <S,G>)

If the requested multicast flow is not part of any list, the join message SHOULD be discarded by the AN. This default behavior can easily be changed by means of a "catch-all" statement in the White list. For instance, adding (<S=*,G=*>) in the White List would make the default behavior to accept join messages for a multicast flow that has no other match on any list.

When the AN receives a "leave" request, it terminates replication of the multicast flow.

If the AN receives a Provisioning message which updates an existing multicast service profile, the AN MUST review the status of active flows on all ports to which the updated profile is currently assigned. Similarly, if a Port Management message assigns a new multicast service profile to a given port, the AN MUST review all active flows on that port. If the most specific match for any flow is a Black list entry, the flow MUST be terminated immediately. If any of the remaining flows do not match an entry in the White list,

they also MUST be terminated immediately. White listed flows MUST be allowed to continue.

6.2.4. Procedures For Conditional Access With Grey Lists

6.2.4.1. Provisioning

The NAS provisions named multicast service profiles containing Grey lists on the AN using the Provisioning message containing one or more Multicast-Service-Profile TLVs. The NAS MAY update the contents of these profiles from time to time as required, by sending additional Provisioning messages with Multicast-Service-Profile TLVs containing incremental modifications to the existing Grey lists or replacements for them.

The NAS assigns a specific multicast service profile to an individual access line using the Port Management message containing a Multicast-Service-Profile-Name TLV. The NAS MAY change profiles on the line by sending a subsequent Port Management message identifying a new profile.

The NAS MAY perform admission control for grey-listed flows. In that case, the NAS MUST NOT include the MRepCtl-CAC TLV in a Provisioning message sent to the AN. Alternatively, the NAS MAY enable admission control at the AN for Grey-listed flows. To do this, it MUST include the MRepCtl-CAC TLV in a Provisioning message sent to the AN and MUST also provide a Bandwidth- Allocation TLV in a Port Management message for each access line.

6.2.4.2. Multicast Service Procedures

The conditional access with Grey lists capability assumes that IGMP/MLD requests are terminated on the AN. When the AN receives a "join" request, it MUST determine whether there is a match to the requested flow in the Grey list of the multicast service profile provisioned against the given access line. If there is no match, the request is discarded. Otherwise, the AN MUST send a Multicast Admission Control message to the NAS with content identifying the access line and the multicast flow to be added. As indicated in Section 4.4, the AN MAY add information identifying the requestor by IP address and/or MAC address.

If the NAS decides to enable the flow, it MUST send a Multicast Replication Control request to the AN to replicate the flow to the access line with the Result field set to Nack (0x1), as described in Section 4.3.1.

When the AN receives the Multicast Replication Control request, it

performs admission control if admission control has been enabled as described in the previous section. If admitting the flow would cause the committed multicast bandwidth at the access line to exceed the provisioned limit, the AN reports an error to the NAS as described in Section 4.3.2. Otherwise it replicates the multicast flow as requested.

If the NAS decides not to permit the flow, it MAY send a Multicast Replication Control message in response to the Multicast Admission Control message to allow the AN to update its internal records. The content of this message is described in Section 4.4.2.

When the AN receives a "leave" request, it MUST terminate replication of the flow to the access line. It MUST then send a Multicast Admission Control message to the NAS indicating the deletion. The NAS updates its internal records but MUST NOT respond to the message.

If the AN receives a Provisioning message which updates an existing multicast service profile, the AN MUST review the status of active flows on all ports to which the updated profile has been assigned. Similarly, if a Port Management message that assigns a new profile to a given port, the AN MUST review all active flows on that port. In either case, if any flow does not match an entry in the Grey list, it MUST be terminated immediately.

6.2.5. Procedures For Delegated Bandwidth

6.2.5.1. Provisioning

The NAS SHOULD provision an initial amount of delegated multicast bandwidth for each access line using the Port Management message containing the Bandwidth-Allocation TLV.

> If it fails to do so and a value has not been provisioned on the AN by other means, the AN will be forced to request a bandwidth allocation as soon as it receives a "join" request.

The NAS MAY at any time force an update of the amount of delegated bandwidth by the same means.

6.2.5.2. Multicast Service Procedures

The delegated bandwidth capability assumes that IGMP/MLD requests are terminated on the AN. When the AN receives a "join" request, it checks whether it has sufficient remaining uncommitted multicast bandwidth on the access line to accommodate the new multicast flow. If not, it MAY send a request to the NAS for an increased allocation of delegated bandwidth, using the Bandwidth Reallocation Request

message. The NAS MUST return a Bandwidth Transfer message indicating whether it has granted the request, and if so, what is the new amount of delegated bandwidth.

If the AN has sufficient uncommitted multicast capacity to admit the request, either originally or as the result of a successful request to the NAS, it replicates the requested flow to the access line. Otherwise it discards the request.

When the AN receives a "leave" request for an active flow, it ceases replication.

The NAS or AN MAY at some point detect that their respective views of the amount of delegated bandwidth are inconsistent. If so, they can recover using procedures described in Section 4.5 and Section 4.6. As a further aid to synchronization, either the NAS or the AN MAY from time to time check the peer's view of the amount of delegated bandwidth using the Delegated Bandwidth Query message.

The NAS or AN MAY at any time release bandwidth to the peer using an autonomous Bandwidth Transfer message. The contents of this message are described in Section 4.6.

6.3. Combinations of Multicast Capabilities

6.3.1. Combination of Conditional Access With White and Black Lists and Conditional Access With Grey Lists

If conditional access with White and Black lists is combined with conditional access with Grey lists, provisioning of the multicast service profiles is as described in Section 6.2.3.1 except that multicast service profiles will also include Grey lists. Admission control is enabled independently on the AN for White lists by including the White-list-CAC TLV in the Provisioning message and for Grey lists by including the MRepCtl-CAC TLV in the Provisioning message. The Bandwidth-Allocation TLV provisions an amount that applies to both White- and Grey- listed flows if admission control is enabled for both.

With regard to multicast service procedures, one point of difference from the individual capabilities must be noted. This is an interaction during the profile matching procedure. The AN MUST seek the best match amongst multiple lists as described in Section 6.2.3.2. However, if there are multiple matches of equal precision, the order of priority is Black list first, Grey list second, and White list last.

Once profile matching has been completed, processing of a "join"

request is as described in Section 6.2.3.2 for White or Black listed flows or Section 6.2.4.2 for Grey listed flows. Requests that do not match any list SHOULD be discarded.

When the AN receives a "leave" request, it MUST terminate replication of the flow to the access line. If the flow was Grey-listed, the AN MUST then send a Multicast Admission Control message to the NAS indicating the deletion. Thus the AN needs to retain the fact that the flow was Grey-listed for the life of the flow.

If the AN receives a Provisioning message which updates an existing multicast service profile, the AN MUST review the status of active flows on all ports to which the updated profile is currently assigned. Similarly, if a Port Management message assigns a new multicast service profile to a given port, the AN MUST review all active flows on that port. If any flow has its most specific match in a Black list entry, it MUST be terminated immediately. If any of the remaining flows do not match an entry in the White or Grey list, they MUST also be terminated immediately. Finally, if any remaining flows were originally admitted because they were White-listed, but after the change they are Grey- listed, the AN MUST generate a Multicast Flow Query response message autonomously as if it were responding to a Multicast Flow Query request, listing all such flows. These flows MUST be allowed to continue until the NAS or the subscriber terminates them. Flows with their most specific match in the White list MUST be allowed to continue.

The autonomously-generated Multicast Flow Query response message MUST be formatted as if it were a successful response to a request containing no Target and no Multicast-Flow TLV, as described in Section 4.9.2, with the exception that the Transaction-Id MUST be set to all zeroes.

6.3.2. Combination of Conditional Access With Delegated Bandwidth

The provisioning and bandwidth management procedures of Section 6.2.5 apply in addition to the procedures in Section 6.2.3, Section 6.2.4, or Section 6.3.1 as applicable. Admission control follows the rules for conditional access in terms of matching flows against White and Black and/or Grey lists and performing or not performing bandwidth checks at the AN, but the amount of bandwidth used by the AN to perform admission control is negotiable as described in Section 6.2.5.2.

6.3.3. Combination of NAS-Initiated Replication with Other Capabilities

NAS-initiated replication can coexist with the other capabilities, but some means must exist to prevent double replication of flows.

The simplest way to do this is to terminate all IGMP/MLD requests on the AN, so that NAS-initiated replication is stimulated only by signalling through other channels. Other arrangements are possible, but need not be discussed here.

Assuming the necessary separation of responsibilities, the only point of interaction between NAS-initiated replication and the other multicast capabilities is in the area of admission control. Specifically, inclusion of the MRepCtl-CAC TLV in a Provisioning message and the Bandwidth-Allocation TLV in a Port Management message enables admission control by the AN for flows added by Multicast Replication Control messages, regardless of whether they are part of NAS-initiated replication or Grey list multicast service processing. Conversely, non inclusion of the MRepCtl-CAC TLV in Provisioning messages to the AN enables admission control by the NAS for flows added by Multicast Replication Control messages, regardless of whether they are part of NAS- initiated replication or Grey list multicast service processing. Admission Control for white flows can also be enabled independently on the AN by inclusion by the NAS of the White-List-CAC TLV in the Provisioning message.

6.3.4. Combinations of Committed Bandwidth Reporting with Other Multicast Capabilities

Committed bandwidth reporting can take independently of which other multicast capabilities have been negotiated. However, some combinations do not make sense because of redundancy. In particular, the NAS obtains the same information that committed bandwidth reporting gives if the only other capabilities operating are NAS-initiated replication and/or conditional access with Grey lists.

7. Security Considerations

The security considerations of ANCP are discussed in [I-D.ietf-ancp-protocol] and in [I-D.ietf-ancp-security-threats]. [Probably need to say more, but will do so later.]

8. IANA Considerations

RFC EDITOR'S NOTE: Please replace XXXX with the RFC number of this document.

This document defines the following additional values within the GSMPv3 Message Type Name Space registry, under the new heading "Multicast Extensions To ANCP/GSMPv3":

Message Name	Message Number	Status	Reference
Multicast Replication Control	144		RFC XXXX
Multicast Admission Control	145		RFC XXXX
Bandwidth Reallocation Request	146		RFC XXXX
Bandwidth Transfer	147		RFC XXXX
Delegated Bandwidth Query	148		RFC XXXX
Multicast Flow Query	149		RFC XXXX
Committed Bandwidth Report	150		RFC XXXX

This document defines the following additional values for the GSMPv3 Failure Response Message Name Space registry:

Value	Failure Response Message Name	Reference
100	Command error.	RFC XXXX
101	Bad flow address.	RFC XXXX
102	Multicast flow does not exist.	RFC XXXX
103	Invalid preferred bandwidth amount.	RFC XXXX
104	Inconsistent views of delegated bandwidth amount.	RFC XXXX
105	Bandwidth request conflict.	RFC XXXX

This document defines the following additional values within the ANCP TLV Type Registry:

TLV Name	Type Code	Reference
Multicast-Service-Profile	0x13	RFC XXXX
Bandwidth-Allocation	0x15	RFC XXXX
Bandwidth-Request	0x16	RFC XXXX
Multicast-Service-Profile-Name	0x18	RFC XXXX
Multicast-Flow	0x19	RFC XXXX
List-Action	0x21	RFC XXXX
Sequence-Number	0x22	RFC XXXX
White-List-CAC	0x24	RFC XXXX
MRepCtl-CAC	0x25	RFC XXXX
Request-Source-IP	0x92	RFC XXXX
Request-Source-MAC	0x93	RFC XXXX
Report-Buffering-Time	0x94	RFC XXXX
Committed-Bandwidth	0x95	RFC XXXX

This document defines the following additional values for the ANCP Command Code registry:

Command Code Directive Name	Command Code Value	Reference
Add	0x01	RFC XXXX
Delete	0x02	RFC XXXX
Delete All	0x03	RFC XXXX
Admission Control Reject	0x04	RFC XXXX
Conditional Access Reject	0x05	RFC XXXX
Admission Control and Conditional Access Reject	0x06	RFC XXXX

This document defines the following additional values for the ANCP Capability registry:

Capability Type Name	Capability Type Code	Reference
NAS-initiated replication	0x3	RFC XXXX
Committed bandwidth reporting	0x5	RFC XXXX
Conditional access with white and black lists	0x6	RFC XXXX
Conditional access with grey lists	0x7	RFC XXXX
Bandwidth delegation	0x8	RFC XXXX

9. Acknowledgements

The authors would like to acknowledge Wojciech Dec for providing useful input to this document, Robert Rennison for his help in shaping the definition of the Multicast-Service-Profile TLV, Shridhar Rao for his comments and suggestions and Aniruddha A for his proposal that formed the base of the Multicast Flow Reporting solution. Philippe Champagne, Sanjay Wadhwa and Stefaan De Cnodder provided substantial contributions on the solution for the NAS initiated multicast control use case. Kristian Poscic provided the committed bandwidth reporting use case.

10. References

10.1. Normative References

[I-D.ietf-ancp-protocol]
Wadhwa, S., Moisand, J., Haag, T., Voigt, N., and T. Taylor, "Protocol for Access Node Control Mechanism in Broadband Networks", draft-ietf-ancp-protocol-15 (work in progress), February 2011.

[IEEE48] IEEE, "http://standards.ieee.org/regauth/oui/tutorials/EUI48.html", 2010.

[IEEE64] IEEE, "http://standards.ieee.org/regauth/oui/tutorials/EUI64.html", 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2710] Deering, S., Fenner, W., and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6", RFC 2710, October 1999.

[RFC3292] Doria, A., Hellstrand, F., Sundell, K., and T. Worster, "General Switch Management Protocol (GSMP) V3", RFC 3292, June 2002.

[RFC3376] Cain, B., Deering, S., Kouvelas, I., Fenner, B., and A. Thyagarajan, "Internet Group Management Protocol, Version 3", RFC 3376, October 2002.

[RFC3810] Vida, R. and L. Costa, "Multicast Listener Discovery Version 2 (MLDv2) for IPv6", RFC 3810, June 2004.

10.2. Informative References

[I-D.ietf-ancp-security-threats]
Moustafa, H., Tschofenig, H., and S. Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", draft-ietf-ancp-security-threats-08 (work in progress), July 2009.

[I-D.morin-mboned-igmpmld-error-feedback]
Morin, T. and B. Haberman, "IGMP/MLD Error Feedback", draft-morin-mboned-igmpmld-error-feedback-02 (work in progress), November 2008.

[PIMreg] IANA, "http://www.iana.org/assignments/pim-parameters/pim-parameters.xhtml", 2005.

[RFC4601] Fenner, B., Handley, M., Holbrook, H., and I. Kouvelas, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", RFC 4601, August 2006.

[RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, November 2008.

[RFC5851] Ooghe, S., Voigt, N., Platnic, M., Haag, T., and S. Wadhwa, "Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks", RFC 5851, May 2010.

Appendix A. Example of Messages and Message Flows

This appendix provides an example in which most of the possible message flows for multicast control are illustrated. This appendix is for informational purposes only. In case of discrepancy with text of the body of this document, the text in the body of the document is to be considered as the normative text.

Assume the following, for a given access port:

- The basic subscribed service is white-listed. The AN will be responsible for admission control for this service.
- Some premium services are available, but requests for these services must be referred to the policy server for proper credit processing. For this reason they are grey-listed. The NAS will be responsible for admission control for these services.
- The subscriber has asked that certain services be blocked so that his children cannot view them. These services are black-listed.
- All of the above services are Source-Specific Multicast (SSM). In addition, by means which bypass the AN, the subscriber can signal intent to join an on- line game service which is Any-Source Multicast (ASM). The NAS is responsible for admission control for this service.
- Bandwidth delegation is in effect to share video bandwidth between the AN and the NAS.

The stated conditions require the use of four of the five capabilities specified in this memo.

A.1. Provisioning Phase

Assume that capability negotiation has been completed between the AN and NAS and that the set of negotiated capabilities includes the following four multicast capabilities: NAS-initiated replication, conditional access with white and black list, conditional access with grey list, and bandwidth delegation. At this point, the NAS can provision the service profiles on the AN and enable admission control at the AN for white-listed flows. To do this, the NAS sends the AN a Provisioning message containing this information. An example message providing the profile for our assumed subscriber is shown in Figure 20. The message has the following contents:

- Message type is 93.

- The Result and Code fields in the header are set to zeroes, as specified in the ANCP base protocol document.

- A Transaction identifier is assigned by the NAS.

- The Multicast-Service-Profile TLV (of which typically there would be multiple instances) contains a Multicast-Service-Profile-Name TLV (with a length of 20 octets assumed for the example) and three List-Action TLVs, one each for the White, Grey, and Black lists within the profile. The White list flows come in two sets of group addresses: 233.252.0.0/29, coming from a server at 192.0.2.15, and 233.252.0.32/29, coming from a server at 192.0.2.16. The Grey listed flows are in the band 233.252.0.64/29, coming from a server at 192.0.2.21. Finally, the Black list flows are two individual flows that happen to overlap with the Grey list band: 233.252.0.65, and 233.252.0.69, also with source 192.0.2.21.

- The White-List-CAC TLV indicates that the AN does admission control on White-listed flows.

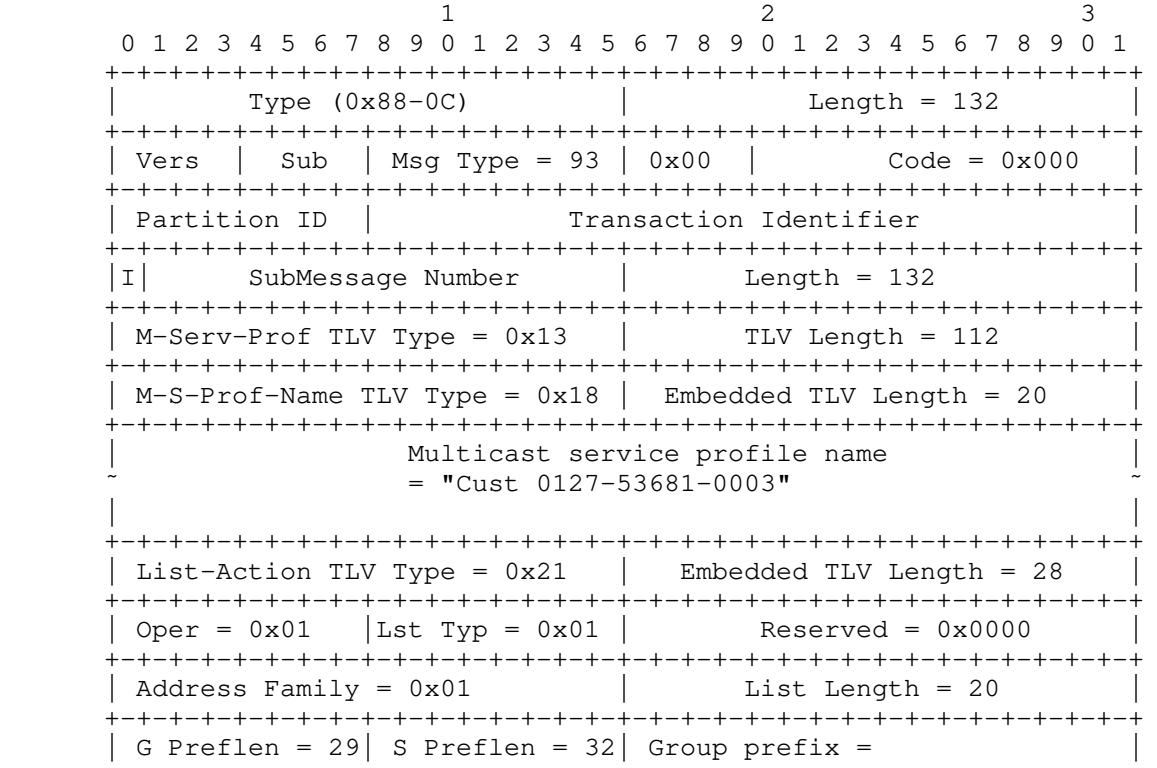

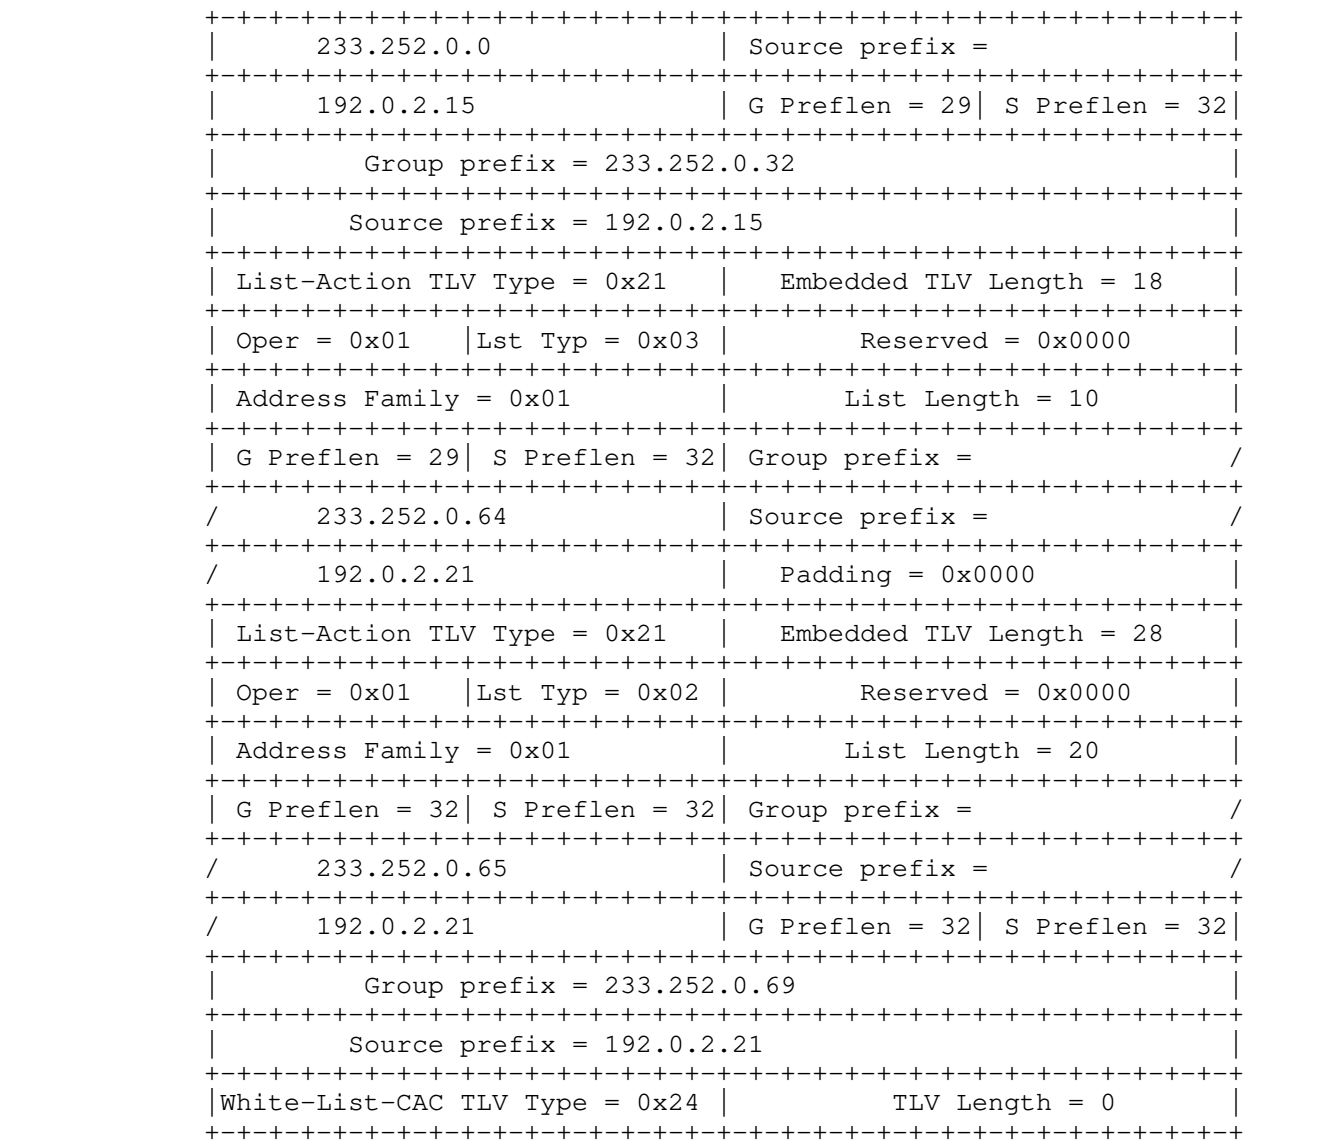

TLV lengths are given in decimal for easier understanding. Note that the padding after the middle List-Action TLV is counted as part of length of the Multicast-Service-Profile TLV, but is not included in the length of that List-Action TLV. Note also that the Length field in the message header, unlike those in the TLVs, includes the message header itself, as required by [RFC3292]. Finally, note that the Provisioning message does not include a MRepCtl-CAC TLV since in our example admission control for Grey listed flows and for NAS-initiated

replication is performed by the NAS.

Figure 20: Example Provisioning Message

As soon as the AN port comes up, the AN sends an ANCP PORT_UP message to the NAS specifying the Access Loop Circuit ID. The NAS replies with an ANCP Port Management message that, together with the other parameters, includes the multicast service profile name to be associated to that port along with the initial amount of delegated bandwidth. The corresponding message flow is illustrated in Figure 21.

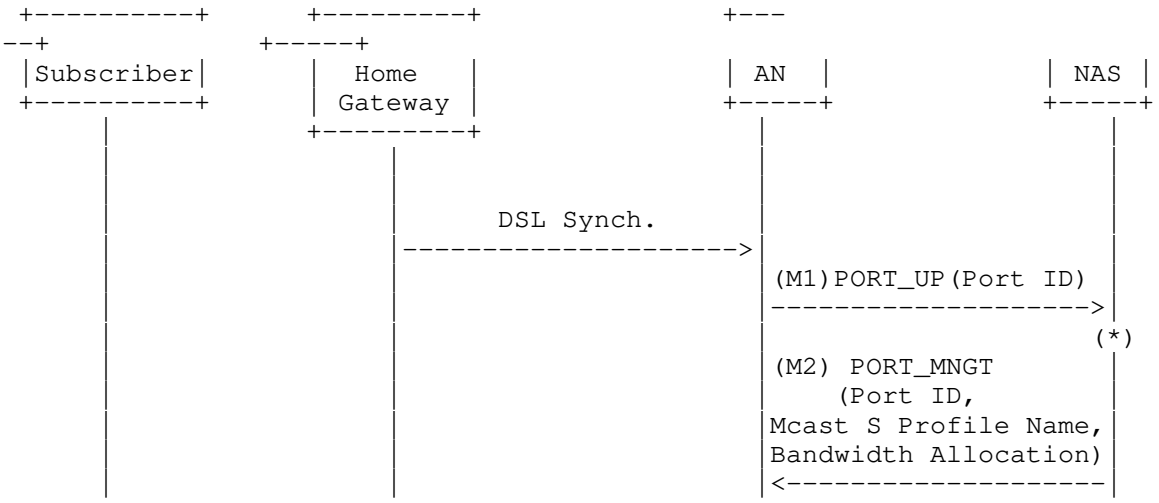

(*) The NAS may optionally seek direction from an external Autorization/Policy Server

Figure 21: Configuring an AN Port With Multicast Service Profile ID and Delegated Bandwidth amount

The Port Management message will typically contain other TLVs but our example (Figure 22) just shows the Target, Multicast- Service-Profile-Name, and Bandwidth-Allocation TLVs. The Target TLV identifies the subscriber line, the Multicast-Service-Profile-Name TLV is identical to the one contained in the Provisioning message, and the Bandwidth-Allocation TLV provides just enough bandwidth (2000 kbits/s) for one channel to start with.

The following fields in the Port Management message header are shown with specific values either as directed by the base protocol document or for the sake of our example:

- Message Type is 32.
- Result is set to Nack (0x01) for this example.
- Code is 0x000.
- Port is set to 0.
- Event Sequence Number, the R flag and the other bits marked x, Duration, the Event Flags, and the Flow Control Flags are all irrelevant for this function and are set to 0.
- Function is set to 0x8, "Configure Connection Service Data".
- X-Function is set to 0.
- Tech Type is 0x05 (DSL).
- Block lengths are calculated assuming a Circuit-Id length of 4 in our example. Recall that the example Multicast-Service-Profile-Name TLV length is 20.

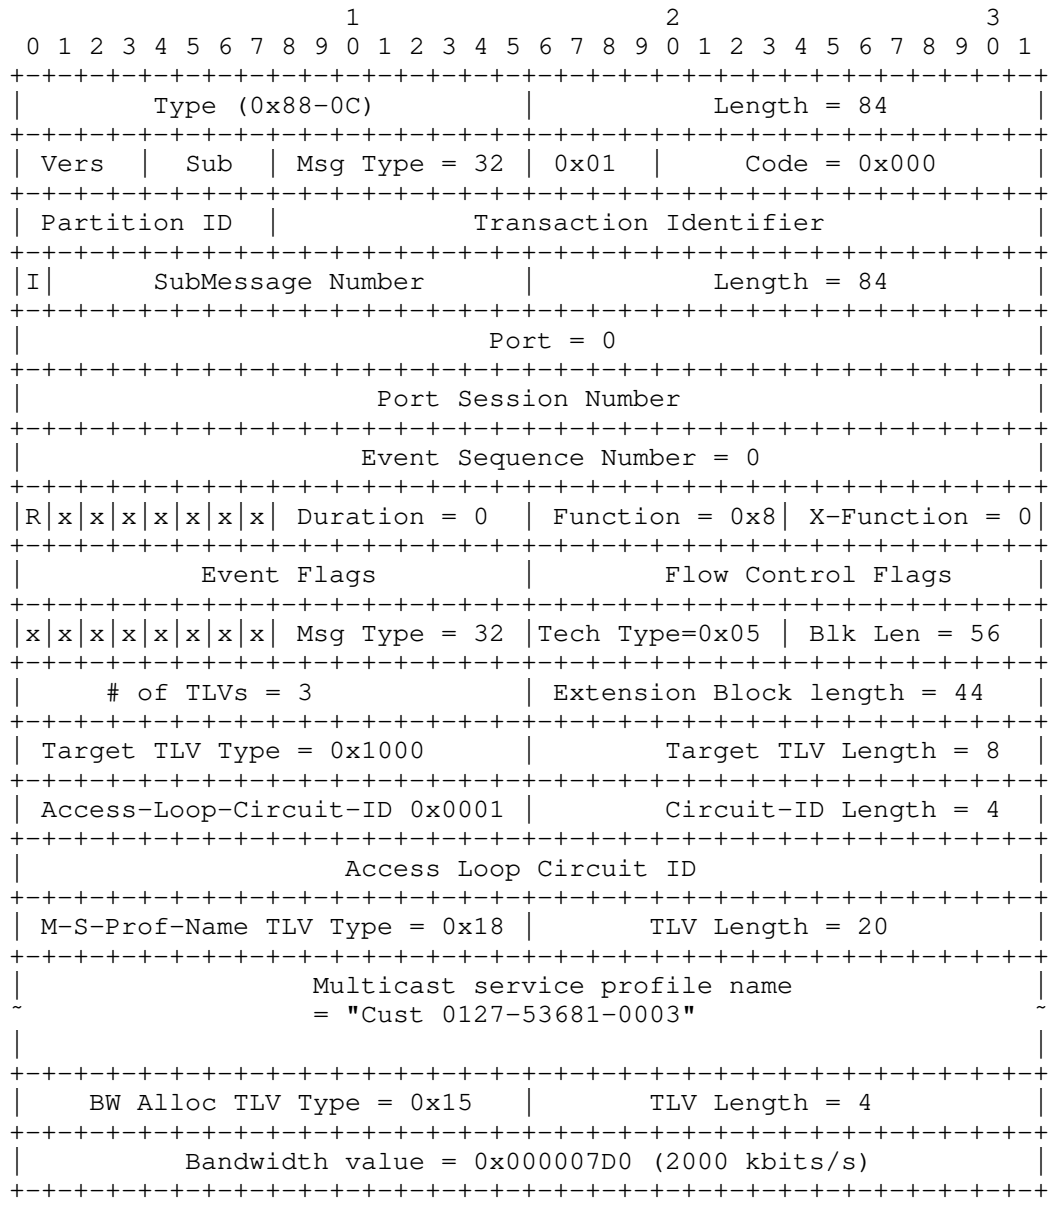

Figure 22: Example Port Management Message

A.2. Handling a Grey-Listed Flow

Suppose now that the subscriber chooses to watch the premium channel characterized by source 192.0.2.21, group 233.252.0.67. Upon receiving the Join request, the AN matches it against the multicast service profile for the port and determines that it is a Grey-listed flow. Figure 23 illustrates the resulting ANCP message flow for the case of a simple join and leave, when admission control for Grey-listed flows is not activated on the AN. To start the flow, the AN sends a Multicast Admission Control request (M1) to the NAS. The NAS decides whether flow can be admitted, applying both policy and bandwidth criteria. It returns its decision (positive in this example) in a Multicast Replication Control message (M2). Later, when the subscriber leaves the flow, the AN informs the NAS by sending another Multicast Admission Control message.

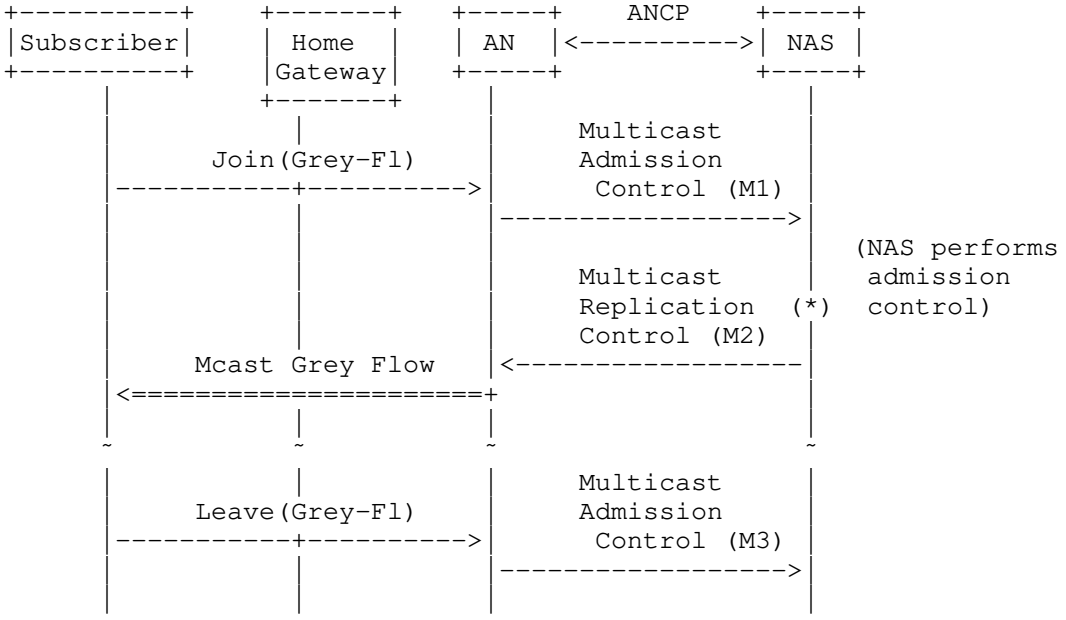

Grey-Fl : Multicast Flow matching an entry in Grey List

(*) The NAS may optionally seek direction from an external Authorization/Policy Server

Figure 23: Successful Join/Leave Operations, Grey- Listed Flow

The Multicast Admission Control message M1 contains:

- an ANCP Header with:
  - Message-Type = 145 - Multicast Admission Control;
  - Result= 0x0 (Ignore);
  - Transaction-ID = Transaction-ID maintained by AN;
- a Target TLV identifying the AN Port
- a Command TLV containing:
  - Command Code = Add (0x01);
  - Accounting = 0;
  - a Multicast-Flow embedded TLV indicating the SSM multicast flow (Flow Type = 0x02) for which the AN received the IGMP Join: IPv4 (0x01) Group address= 233.252.0.67, IPv4 (0x01) Source Address = 192.0.2.21;
  - a Request-Source-IP embedded TLV containing the IGMP join source IP (192.0.2.100).

The Multicast Admission Control message M1 is illustrated below:

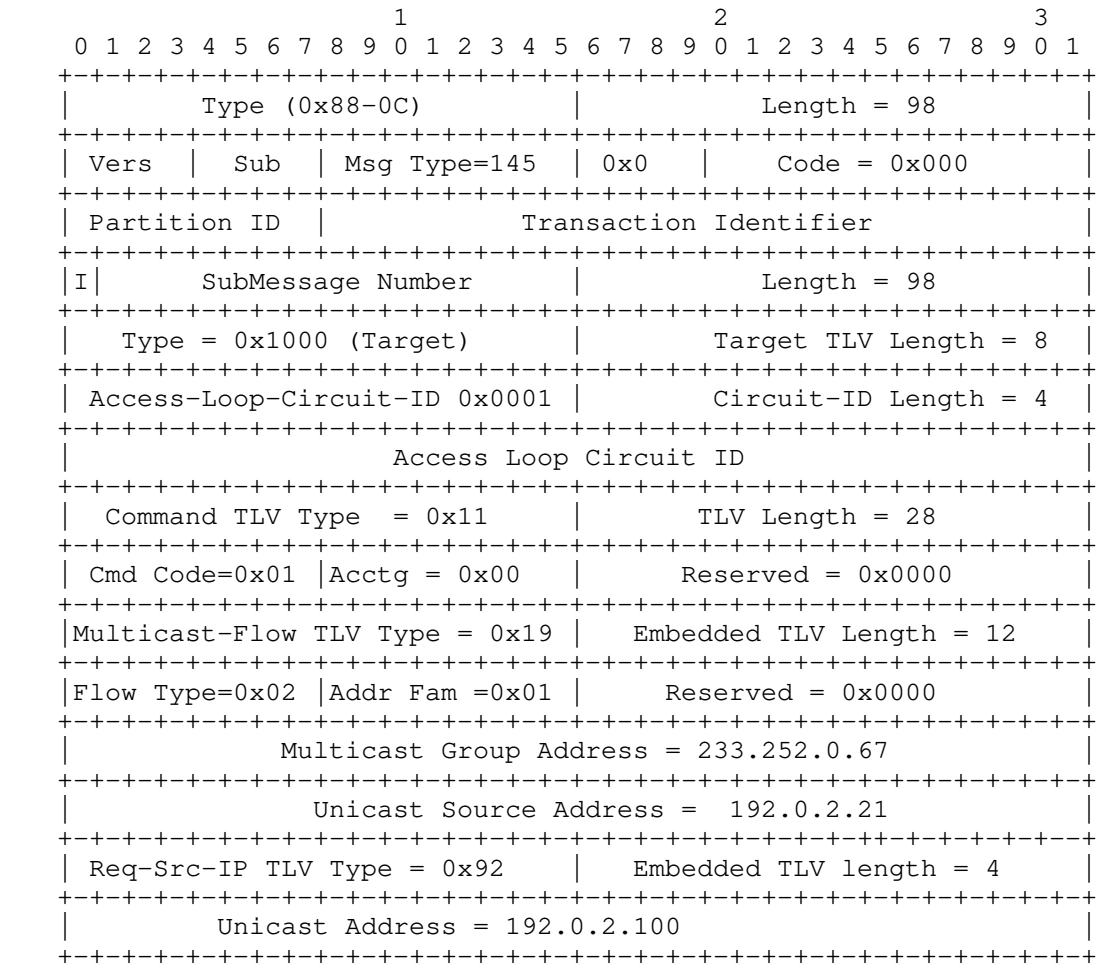

Multicast Admission Control Message Seeking To Add A Flow

The Multicast Replication Control message M2 contains:

- an ANCP Header with:
  - Message Type = 144 - Multicast Replication Control;
  - Result= 0x1 (NAck);
  - Transaction-ID = Transaction-ID maintained by NAS;

- a Target TLV identifying the AN Port;
- a Command TLV containing:
  - Command Code = Add (0x01);
  - Accounting = 1 (begin flow accounting), since in our example the operator wants accounting on this flow.
  - a Multicast-Flow embedded TLV indicating the SSM multicast flow (Flow Type = 0x02) that the NAS is admitting for this access port: IPv4 (0x01) Group address= 233.252.0.67, IPv4 (0x01) Source Address = 192.0.2.21.

The Multicast Admission Control message M2 is illustrated below.

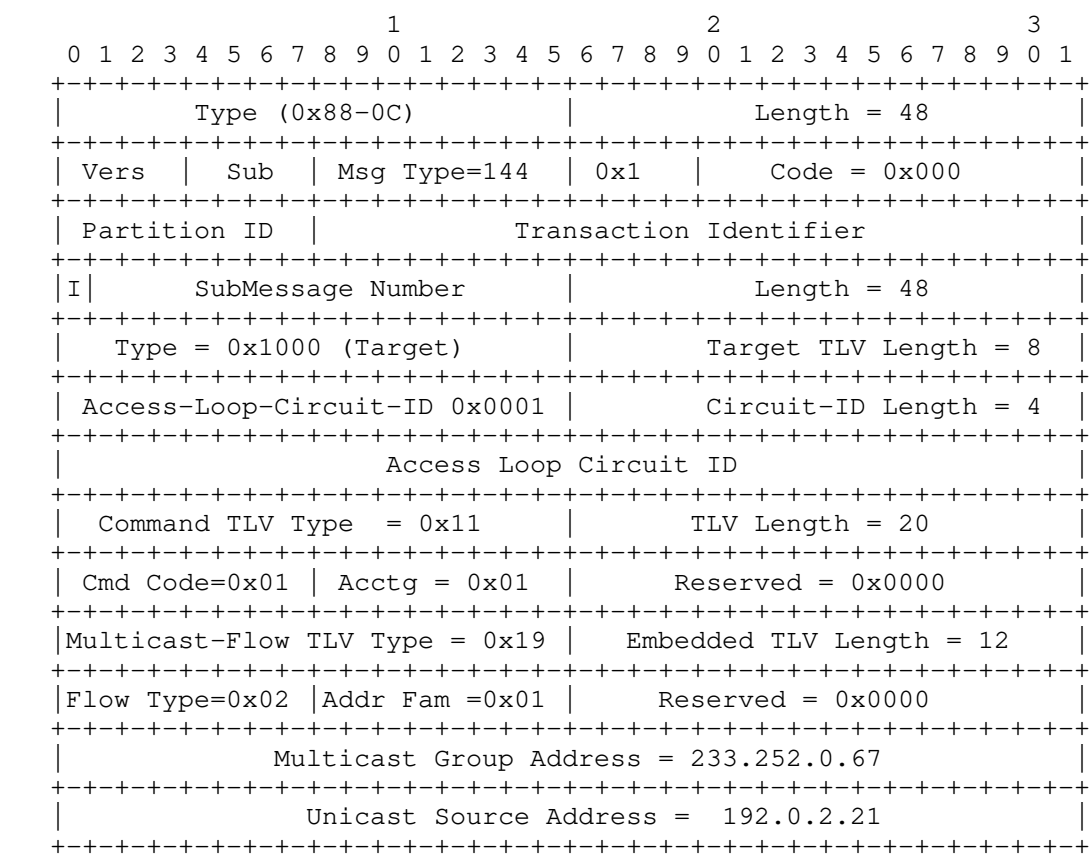

Figure 24: Multicast Replication Control Message Admitting A Flow

The Multicast Admission Control message M3 advising the NAS that the flow has been terminated contains:

- an ANCP Header with:
  - Message-Type = 145 - Multicast Admission Control
  - Result= 0x0 (Ignore)
  - Transaction-ID = Transaction-ID maintained by AN
- a Target TLV identifying the AN Port
- a Command TLV containing:
  - a Command Code = Delete (0x02);
  - Accounting = 0;
  - a Multicast-Flow embedded TLV indicating the SSM multicast flow (Flow Type = 0x02) for which the AN received the IGMP leave: IPv4 (0x01) Group address= 233.252.0.67, IPv4 (0x01) Source Address = 192.0.2.21.
  - a Request-Source-IP embedded TLV containing the IGMP leave request source, IPv4 (0x01) address 192.0.2.100.

The Multicast Admission Control message M3 is illustrated below.

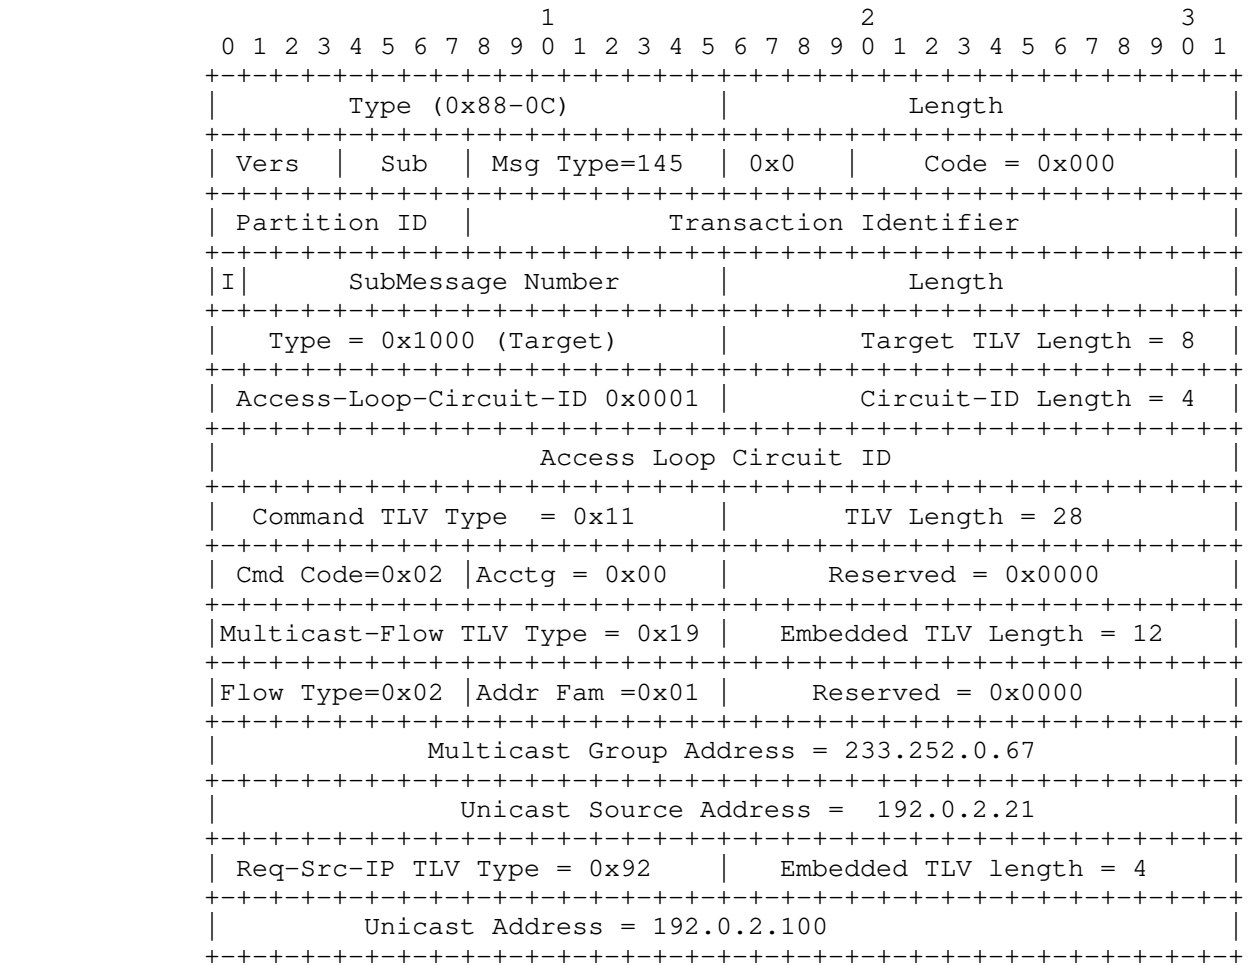

Multicast Admission Control Message Signalling flow Termination

A.3. Handling White-Listed Flows

The NAS has enabled White list admission control on the AN, and the bandwidth delegation capability has been negotiated. White listed flows in themselves require no messages to the NAS, either upon admission or upon termination, but the AN may request an increase in the amount of delegated bandwidth if it needs the increase to admit a flow.

Consider an example where the AN has already admitted one White-listed flow, thereby using up the initially provisioned amount of delegated bandwidth (2000 kbits/s). A request is received to join a

new flow in the White list range. The AN chooses to send a Bandwidth Reallocation Request message to the NAS, requesting that the delegated bandwidth allocation be increased to 4000 kbits/s at a minimum, and preferably to 6000 kbits/s.

In our example, the NAS is managing bandwidth tightly, as witnessed by its minimal initial allocation of just enough for one flow. It is willing to provide the minimum additional amount only, and therefore returns a Bandwidth Transfer message where the delegated bandwidth value is given as 4000 kbits/s. With this amount, the AN is able to admit the second White-listed flow. The AN could send a similar Bandwidth Transfer message back to the NAS bringing the delegated bandwidth amount back down to 2000 kbits/s when one of the flows is terminated, but this shows nothing new and is omitted.

As one more point of illustration, suppose that the NAS chooses to audit the current amount of delegated bandwidth to ensure it is synchronized with the AN. It sends a Delegated Bandwidth Query request message to the AN, and receives a Delegated Bandwidth Query response message with the current allocation as the AN sees it.

The complete message flow is shown in Figure 25.

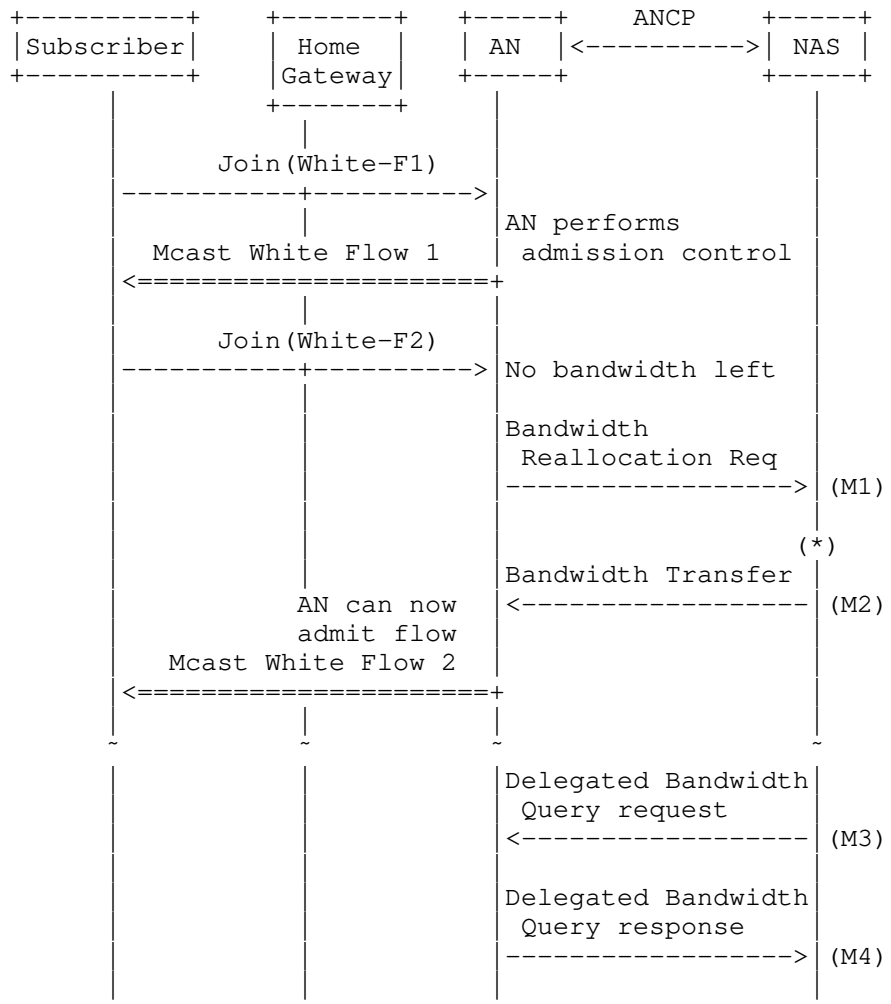

(*) The NAS may optionally seek direction from an external Authorization/Policy Server

Figure 25: Successful Join/Leave Operations, Grey-Listed Flow

The Bandwidth Reallocation Request message (M1) is shown in Figure 26. The contents require little explanation. The Message Type for the Bandwidth Reallocation Request is 146. The Result field is set to 0x0 (Ignore). Besides the Target, the message has one other TLV, the Bandwidth- Request, with a TLV Type of 0x16. The TLV contains Required Amount and Preferred Amount fields, set to 4000 and 6000 kbits/s respectively.

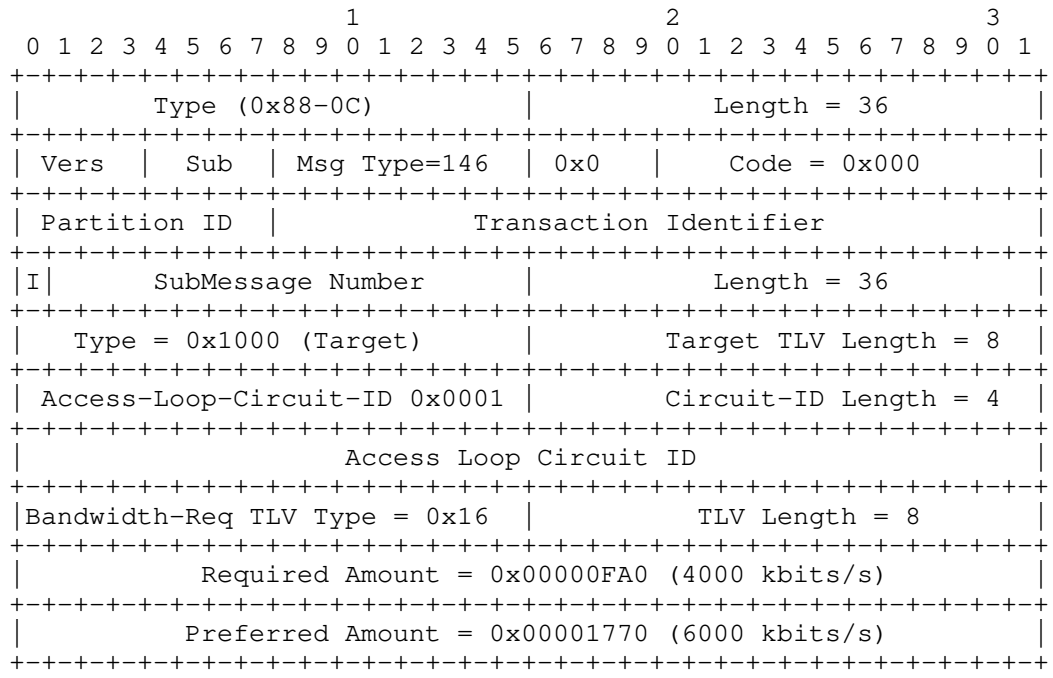

Figure 26: Bandwidth Reallocation Request Message

The Bandwidth Transfer message (M2) is shown in Figure 27. Again, the contents are easily understood. The Message Type for the Bandwidth Transfer message is 147. The Result field is set to Success (0x3). The message contains the Target TLV and the Bandwidth- Allocation TLV. The latter has a TLV Type of 0x15 and contains a Delegated Amount field, set to 4000 kbits/s.

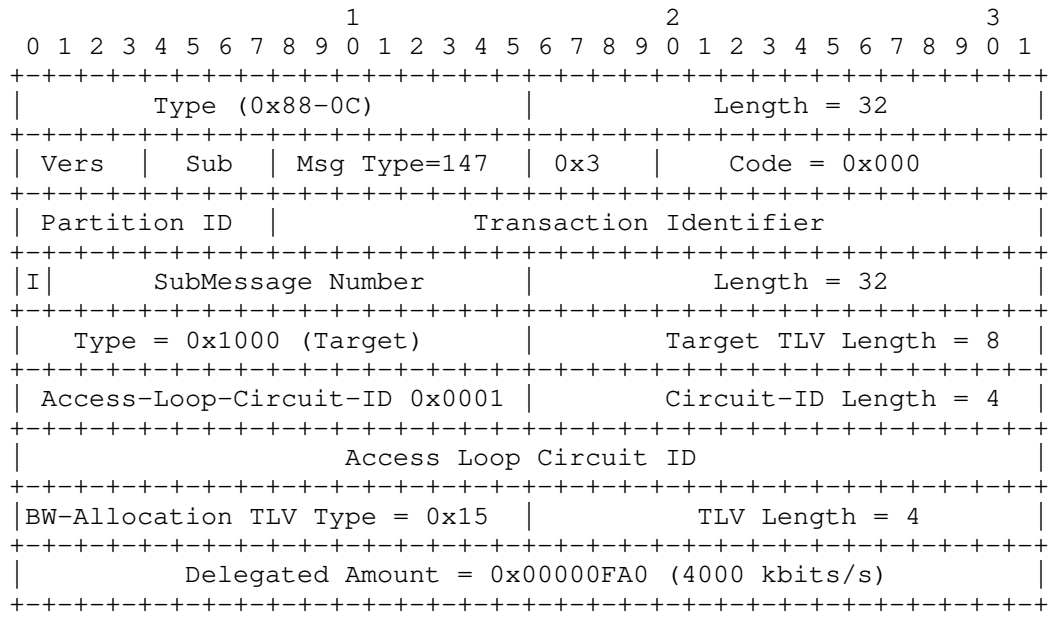

Figure 27: NAS Response, Bandwidth Transfer Message

The Delegated Bandwidth Query request message (M3) is shown in Figure 28. The Message Type for the Delegated Bandwidth Query request message is 148. The Result field is set to AckAll (0x2). The message contains the Target TLV only.

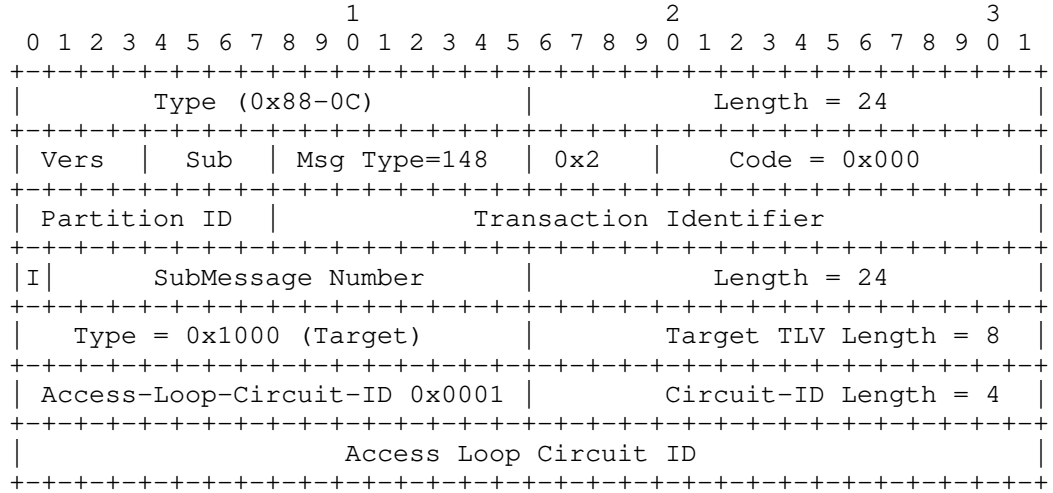

Figure 28: Delegated Bandwidth Query Request Message

Finally, the Delegated Bandwidth Query response message (M4) is shown in Figure 29. The Message Type for the Delegated Bandwidth Query response message is 148. The Result field is set to Success (0x3). The message contains the Target TLV and the Bandwidth-Allocation TLV with the Delegated Amount field set to 4000 kbits/s.

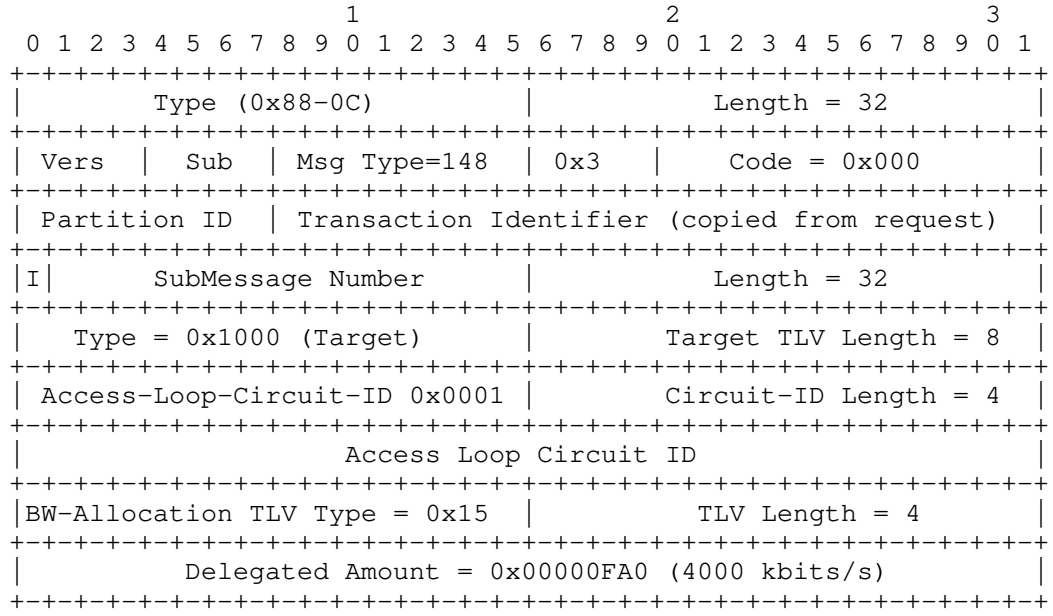

Figure 29: Delegated Bandwidth Query Response Message

A.4. Handling Of Black-Listed Join Requests

This section introduces no new messages, since requests for flows in the Black list are simply ignored. The one thing to point out is the overlap in our example between the set of flows in the Grey list and the flows in the Black list. This does not create any ambiguity, since not only does the Black list have priority for equally good matches, but also the Black list entries are more specific (group prefix lengths of 32 versus 29 in the Grey list) than the Grey list flow prefixes.

A.5. Handling Of Requests To Join and Leave the On-Line Game

The final class of multicast control actions in our example allows the subscriber to enter and leave the on-line game. As described at the beginning of this example, the game uses Any Source Multicast (ASM). Subscriber signalling bypasses the AN, going directly to the NAS (e.g., through a web interface).

When the subscriber requests to join the game, the NAS (after applying policy and bandwidth checks) sends a Multicast Replication Control message to the AN to enable the flow on the port concerned. The AN knows not to apply admission control, since it has not received an MRepCtl-CAC TLV in the Provisioning message. When the subscriber leaves, the NAS sends another Multicast Replication Control message to delete the flow. This message sequence is shown in Figure 30.

It is possible that the NAS finds that there is not enough bandwidth available to accommodate the subscriber's request. In this case, the NAS could send a Bandwidth Reallocation Request message to the AN, asking it to release some of the bandwidth delegated to it. This is not shown in the present example, since the messages are the same as those already presented with the exception that the Preferred Amount in the request will be *less than* or equal to the Required amount, rather than *greater than* or equal to it.

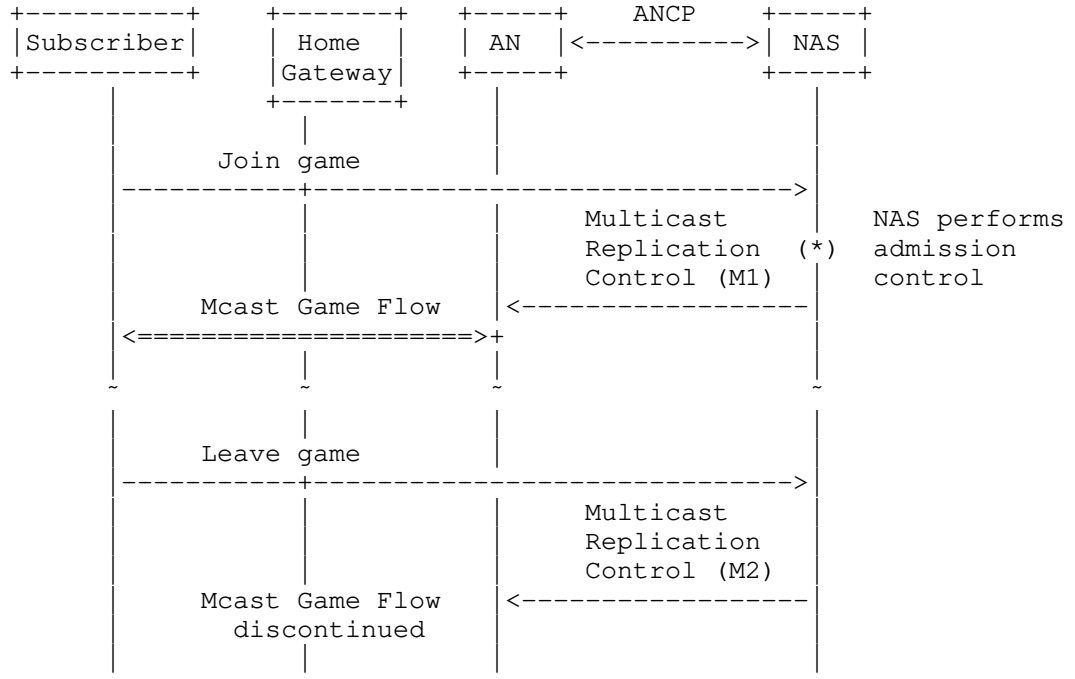

(*) The NAS may optionally seek direction from an external Authorization/Policy Server

Figure 30: NAS-Initiated Flows For On-Line Gaming

Multicast Replication Control message (M1) in Figure 31 looks like

the message in Figure 24 with two exceptions. The first is that the NAS has the option to set the Result field to AckAll (0x02) if it needs positive reassurance that the flow has been enabled. This was not done here to save having to depict a response differing only in the Result field. The larger difference in this example is that the flow description in the Multicast-Flow embedded TLV is that of an ASM multicast group (Flow Type = 0x01) with IPv4 (0x01) group address 233.252.1.100.

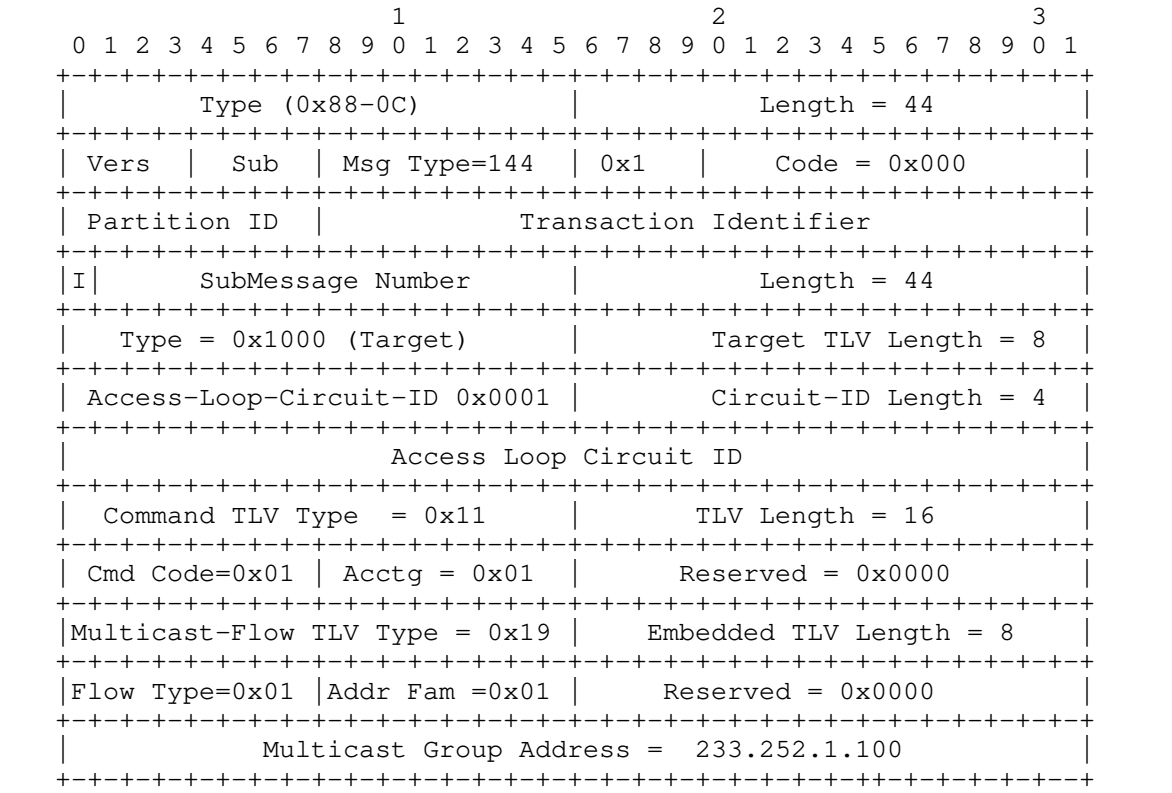

Figure 31: Enabling The Subscriber To Join An On- Line Game

Message M2 terminating the flow when the subscriber leaves the game looks the same as the message in Figure 31 with two exceptions: the Command Code becomes Delete (0x02), and Accounting is set to 0 to turn off flow accounting. Of course, the Transaction Identifier values will differ between the two messages.

A.6. Example Flow For Multicast Flow Reporting

The example in this section is independent of the example in the preceding sections.

Figure 32 illustrates a message flow in a case where the NAS queries the AN about which multicast flows are active on port 10, on port 20 and on port 11 of the AN.

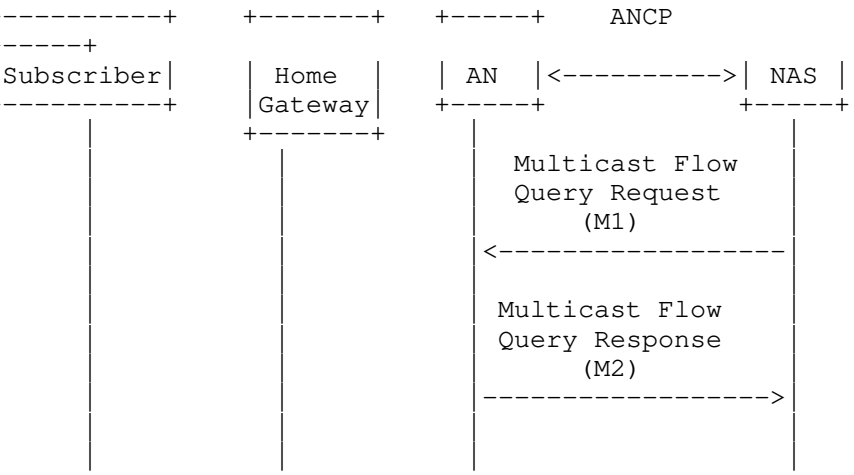

Figure 32: Per Port Multicast Flow Reporting

The Multicast Flow Query Request message (M1) is illustrated in Figure 33. The Message Type is 149. The Result field is set to AckAll (0x2). Three Target TLVs are present, identifying port 10, port 20, and port 11 respectively.

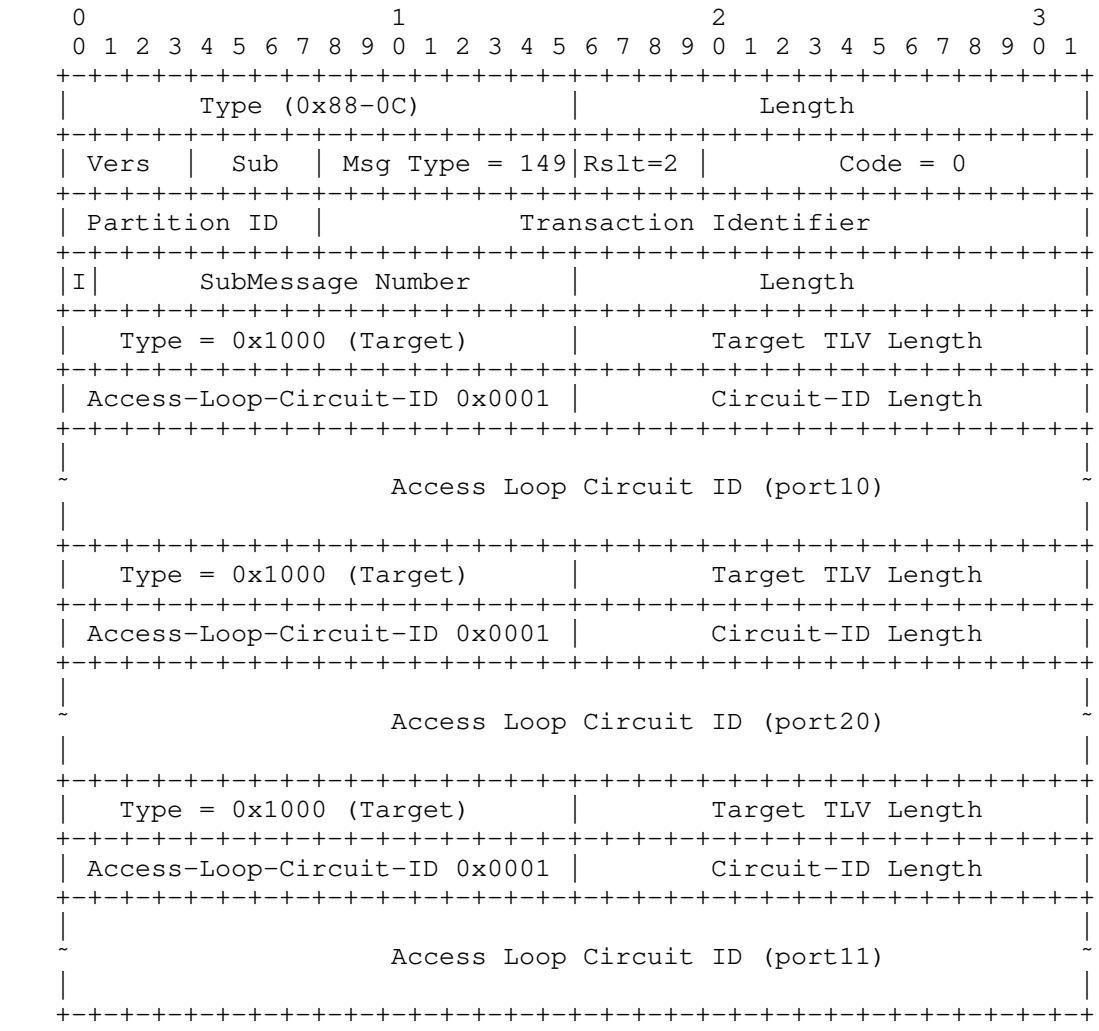

Figure 33: Multicast Flow Query Request Message For Per-Port Multicast Flow Reporting

The Multicast Flow Query Response message (M2) is illustrated in Figure 34. It indicates that there is one active multicast flow [(192.0.2.1, 233.252.2.4)] on port 10, no active multicast flow on port 20 and two active multicast flows [(192.0.2.1, 233.252.2.4) and (192.0.2.2, 233.252.2.10)] on port 11.

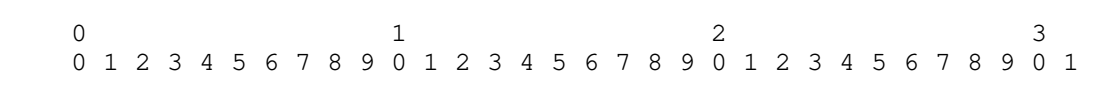

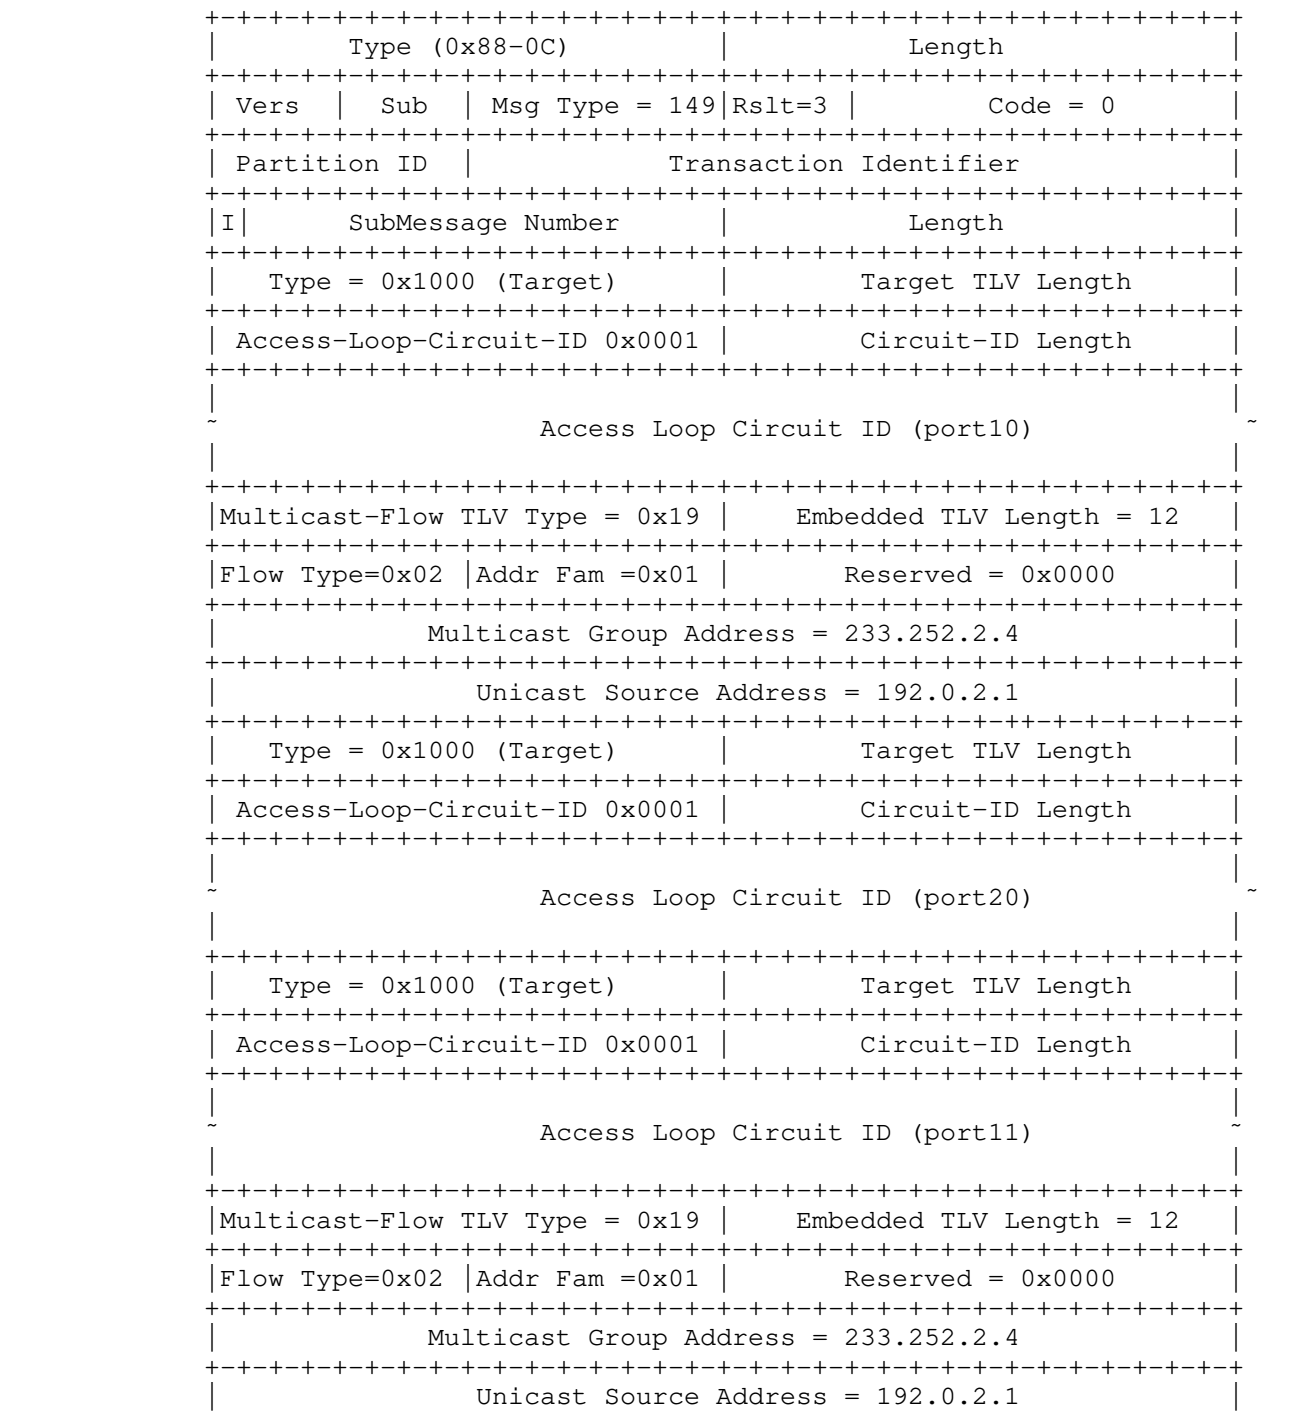
Type (0x88-0C)
Length
Vers
Sub
Msg Type = 149
Rslt=3
Code = 0
Partition ID
Transaction Identifier
I
SubMessage Number
Length
Type = 0x1000 (Target)
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port10)
Multicast-Flow TLV Type = 0x19
Embedded TLV Length = 12
Flow Type=0x02
Addr Fam =0x01
Reserved = 0x0000
Multicast Group Address = 233.252.2.4
Unicast Source Address = 192.0.2.1
Type = 0x1000 (Target)
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port20)
Type = 0x1000 (Target)
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port11)
Multicast-Flow TLV Type = 0x19
Embedded TLV Length = 12
Flow Type=0x02
Addr Fam =0x01
Reserved = 0x0000
Multicast Group Address = 233.252.2.4
Unicast Source Address = 192.0.2.1

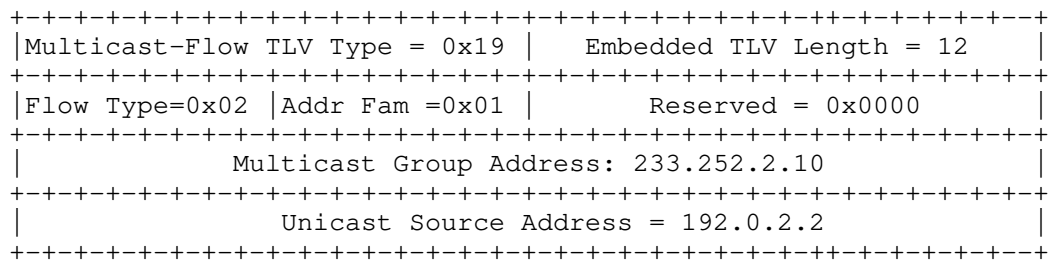

Figure 34: Multicast Flow Query Response message for per-port Mulicast Flow Reporting

Authors' Addresses

Francois Le Faucheur
Cisco Systems
Greenside, 400 Avenue de Roumanille
Sophia Antipolis 06410
France

Phone: +33 4 97 23 26 19
Email: flefauch@cisco.com

Roberta Maglione
Telecom Italia
Via Reiss Romoli 274
Torino 10148
Italy

Phone:
Email: roberta.maglione@telecomitalia.it

Tom Taylor
Huawei Technologies
Ottawa
Canada

Email: tom111.taylor@bell.net

Network Working Group
Internet Draft
Category: Informational
Expiration Date: April 18, 2011

Nabil Bitar
Verizon

Sanjay Wadhwa
Juniper Networks

October 18, 2010

Applicability of Access Node Control Mechanism to PON based Broadband Networks

draft-ietf-ancp-pon-00.txt

Abstract

The purpose of this document is to provide applicability of Access Node Control Mechanism, as described in [ANCP-FRAMEWORK], to PON based broadband access. The need for an Access Node Control Mechanism between a Network Access Server (NAS) and an Access Node Complex (a combination of Optical Line Termination (OLT) and Optical Network Termination (ONT) elements), is described in a multi-service reference architecture in order to perform QoS-related, service-related and Subscriber-related operations. The Access Node Control Mechanism is also extended for interaction between components of the Access Node Complex (OLT and ONT). The Access Node Control mechanism will ensure that the transmission of the information does not need to go through distinct element managers but rather uses a direct device-device communication. This allows for performing access link related operations within those network elements to meet performance objectives.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/1id-abstracts.html.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on April 18, 2011.

Table of Contents

1 Specification Requirements

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

2 Introduction

Passive Optical Networks (PONs) based on BPON and GPON are being deployed across carrier networks. There are two models for PON deployment: Fiber to the curb (FTTC), and Fiber to the Premise (FTTP). In the FTTC deployment, the last mile connectivity is provided over the local loop using Very High Speed DSL. In the FTTP case, PON extends to the premise. In addition, there are three main PON technologies: (1) Broadband PON (BPON), (2) Gigabit PON (GPON), and (3) Ethernet PON (EPON). The focus in the document will be on BPON and GPON in the context of FTTP deployment.

BPON and GPON in FTTP deployments provide large bandwidth in the first mile, bandwidth that is an order of magnitude larger than that provided by xDSL. In the downstream direction BPON provides 622 Mbps per PON while GPON provides 2.4 Gbps. In residential deployments, the number of homes sharing the same PON is limited by the technology and the network engineering rules. Typical deployments have 32 homes per PON.

The motive behind BPON and GPON deployment is providing triple-play services over IP: voice, video and data. Voice is generally low bandwidth but has low-delay, low-jitter, and low packet-loss requirements. Data services (e.g., Internet services) often require high throughput and can tolerate medium latency. Data services may include multimedia content download such as video. However, in that case, the video content is not required to be real-time and/or it is low quality video. Video services, on the other hand, are targeted to deliver Standard Definition or High Definition video content in real-time or near-real time, depending on the service model. Standard Definition content using MPEG2 encoding requires on the order of 3.75 Mbps per stream while High definition content using MPEG2 encoding requires on the order of 15-19 Mbps depending on the level of compression used. Video services require low-jitter and low-packet loss with low start-time latency. There are two types of video services: on demand and broadcast (known also as liner programming content). While linear programming content can be provided over Layer1 on the PON, the focus in this document is on delivering linear programming content over IP to the home, using IP multicast. Video on demand is also considered for delivery over IP using a unicast session model.

Providing simultaneous triple-play services over IP with unicast video and multicast video, VoIP and data requires an architecture that preserves the quality of service of each service. Fundamental to this architecture is ensuring the video content (unicast and multicast) delivered to the user does not exceed the bandwidth allocated to the user for video services. Architecture models often ensure that data is guaranteed a minimum bandwidth and that VoIP is guaranteed its own bandwidth. In addition, QoS control across services is often performed at a Network Access Server (NAS), often referred to as Broadband Network Gateway (BNG) for subscriber management, per subscriber and shared link resources. Efficient multicast video services require enabling multicast services in the access network between the subscriber and the subscriber management platform. In the FTTP PON environment, this implies enabling IP multicast on the Access Node (AN) complex composed of the ONT and OLT, as applicable. This is as opposed to DSL deployments where multicast is enabled on the DSLAM only. The focus in this document will be on the ANCP requirements needed for coordinated admission control of unicast and multicast video in FTTP PON environments between the AN complex and the NAS, specifically focusing on bandwidth dedicated for multicast and shared bandwidth between multicast and unicast.

[ANCP-FRAMEWORK] provides the framework and requirements for coordinated admission control between a NAS and an AN with special focus on DSL deployments. This document proposes the extension of that framework and the related requirements to explicitly address BPON and GPON deployments.

2.1 Terminology

- PON (Passive Optical Network): a point-to-multipoint fiber to the premises network architecture in which unpowered splitters are used to enable the splitting of an optical signal from a central office on a single optical fiber to multiple premises. Up to 32-128 may be supported on the same PON. A PON configuration consists of an Optical Line Termination (OLT) at the Service Provider's CO and a number of Optical Network Units or Terminals (ONU/ONT) near end users, with an optical distribution network (ODN) composed of fibers and splitters between them. A PON configuration reduces the amount of fiber and CO equipment required compared with point to point architectures.

- Access Node Complex (ANX): The Access Node is decomposed by two geographical functions, performed by OLT and ONU/ONT. The general term Access Node (ANX) will be used when describing a functionality which does not depend on the physical location but rather on the "black box" behaviour of OLT and ONU/ONT.

- Optical Line Terminal (OLT): is located in the Service provider's central office. It terminates and aggregates multiple PONs (providing fiber access to multiple premises or neighborhoods) on the user side, and interfaces with the service element (NAS) providing subscriber management.

- Optical Network Terminal (ONT): terminates PON on the network side and provides PON adaptation. The user side interface and the location of the ONT is dictated by the type of network deployment. For a Fiber-to-the-Premise (FTTP) deployment (with Fiber all the way to the apartment or living unit), ONT has Ethernet (FE/GE/MoCA) connectivity with the Home Gateway (HGW)/Customer Premise Equipment (CPE). In case of an MDU (multi-dwelling or multi-tenant unit), a multi-subscriber ONU typically resides in the basement or a wiring closet, and has FE/GE/Ethernet over VDSL connectivity with each CPE. In the case where fiber is terminated outside the premise (neighborhood or curb side) on an ONT/ONU, the last-leg-premise connections could be via existing or new Copper, with xDSL physical layer (typically VDSL). In this case, the Access Node (OLT & ONT together) effectively is a "PON fed DSLAM".

- Network Access Server (NAS): Network element which aggregates subscriber traffic from a number of ANs or ANXs. The NAS is often an injection point for policy management and IP QoS in the access network. It is also referred to as Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS).

- Home Gateway (HGW): Network element that connects subscriber devices to the AN or ANX and the access network. In case of DSL, the Home Gateway is a DSL network termination that could either operate as a Layer 2 bridge or as a Layer 3 router. In the latter case, such a device is also referred to as a Routing Gateway (RG). In the case of PON, it is often a Layer3 routing device with the ONT performing PON termination.

- PON-Customer-ID: This is an identifier which uniquely identifies the ANX and the access loop logical port on the ANX to the customer premise, and is used in any interaction between NAS and ANX that relates to access-loops. Logically it is composed of information containing identification of the OLT (the OLT may be physically directly connected to the NAS), the PON port on the OLT, the ONT, and the port on the ONT connecting to the customer HGW. When acting as a DHCP relay agent, the OLT can encode PON-Customer-ID in the "Agent-Circuit-Identifier" Sub-option in Option-82 of the DHCP messages.

3 Reference Architecture for PON Based Broadband Access Network

The reference architecture used in this document is based on Ethernet aggregation for both of BPON and GPON. Specifically, the following cases are addressed:

- o BPON with Ethernet uplink to the BNG and ATM on the PON side.
- o GPON with Ethernet uplink to the BNG and Ethernet on the PON side.

In case of an Ethernet aggregation network that supports new QoS-enabled IP services (including Ethernet multicast replication), the architecture builds on the reference architecture specified in DSL Forum [TR-101]. The Ethernet aggregation network between a NAS and an OLT may be degenerated to one or more direct physical Ethernet links.

Given the industry's move towards Ethernet as the new access and aggregation technology for triple play services, the primary focus throughout this document is on GPON and BPON with Ethernet between the BNG and the OLT. Figures 1 and 2 depict an end-to-end broadband network with PON access.

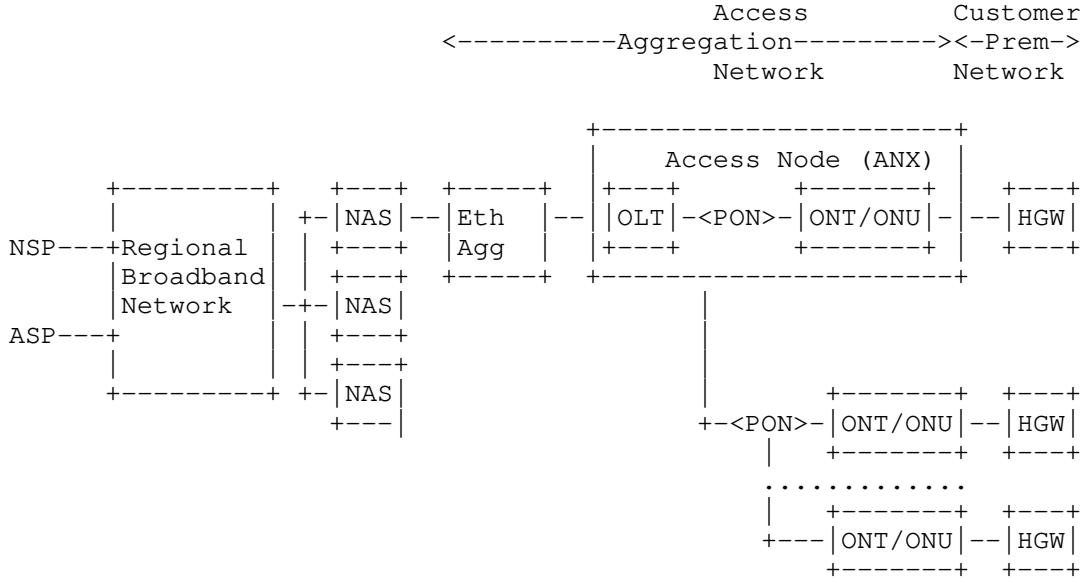

Figure 1. Access Network with PON

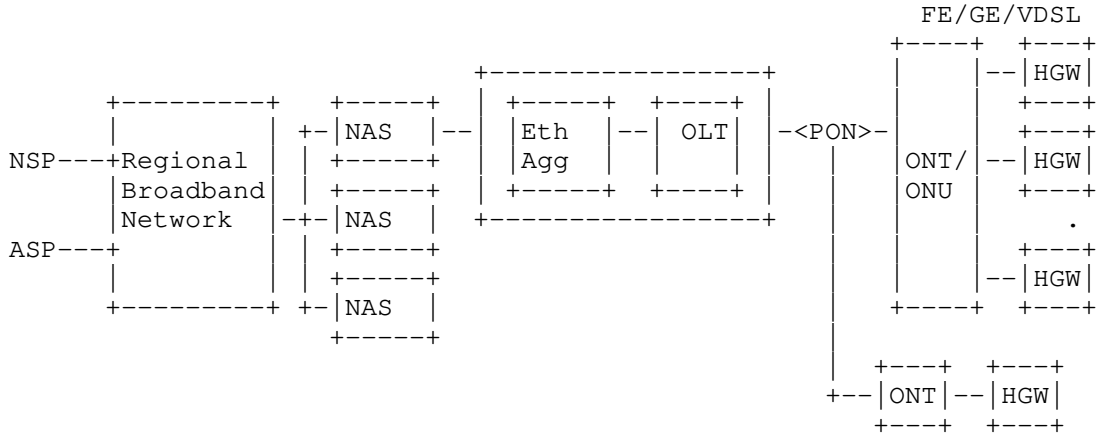

Figure 2. FTTP/FTTC with multi-subscriber ONU serving MTUs/MDUs

3.1 Home Gateway

The Home Gateway (HGW) connects the different Customer Premises Equipment (CPE) to the ANX and the access network. In case of PON, the HGW is a layer 3 router. In this case, the HGW performs DHCP assignment to devices within the home, and performs Network Address and Port Translation (NAPT) between the LAN and WAN side. In case of FTTP, the HGW connects to the ONT over an Ethernet interface. That Ethernet interface could be a physical port or over another medium. In case of FTTP, it is possible to have a single box GPON CPE solution, where the ONT encompasses the HGW functionality as well as the GPON adaptation function.

3.2 PON Access

PON access is composed of the ONT and OLT. PON ensures physical connectivity between the ONT at the customer premises and the OLT. PON framing can be BPON (in case of BPON) or GPON (in case of GPON). The protocol encapsulation on BPON is based on multi-protocol encapsulation over AAL5, defined in [RFC2684]. This covers PPP over Ethernet (PPPoE, defined in [RFC2516]), or bridged IP (IPoE). The protocol encapsulation on GPON is always IPoE. In all cases, the connection between the AN (OLT) and the NAS (BNG) is assumed to be Ethernet in this document.

3.3 Access Node Complex

This is composed of OLT and ONT and is defined in section 2.1.

3.4 Access Node Complex Uplink to the BNG

The ANX uplink connects the OLT to the NAS. The fundamental requirements for the ANX uplink are to provide traffic aggregation, Class of Service distinction and customer separation and traceability. This can be achieved using an ATM or an Ethernet based

technology. The focus in this document is on Ethernet as stated earlier.

3.5 Aggregation Network

The aggregation network provides traffic aggregation towards the NAS. The Aggregation network is assumed to be Ethernet in this document.

3.6 Network Access Server

The NAS is a network device which aggregates multiplexed Subscriber traffic from a number of ANXs. The NAS plays a central role in per-subscriber policy enforcement and QoS. It is often referred to as a Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS). A detailed definition of the NAS is given in [RFC2881]. The NAS interfaces to the aggregation network by means of 802.1Q or 802.1 Q-in-Q Ethernet interfaces, and towards the Regional Network by means of transport interfaces (e.g. GigE, PPP over SONET). The NAS functionality corresponds to the BNG functionality described in DSL Forum TR-101. In addition to this, the NAS supports the Access Node Control functionality defined for the respective use cases in this document.

3.7 Regional Network

The Regional Network connects one or more NAS and associated Access Networks to Network Service Providers (NSPs) and Application Service Providers (ASPs). The NSP authenticates access and provides and manages the IP address to Subscribers. It is responsible for overall service assurance and includes Internet Service Providers (ISPs).The ASP provides application services to the application Subscriber (gaming, video, content on demand, IP telephony etc.). The NAS can be part of the NSP network. Similarly, the NSP can be the ASP.

4 Motivation for explicit extension of ANCP to FTTP PON

The fundamental difference between PON and DSL is that a PON is an optical broadcast network by definition. That is, at the PON level, every ONT on the same PON sees the same signal. However, the ONT filters only those PON frames addressed to it. Encryption is used on the PON to prevent eavesdropping.

The broadcast PON capability is very suitable to delivering multicast content to connected premises, maximizing bandwidth usage efficiency on the PON. Similar to DSL deployments, enabling multicast on the Access Node Complex (ANX) provides for bandwidth use efficiency on the path between the Access Node and the NAS as well as improves the scalability of the NAS by reducing the amount of multicast traffic being replicated at the NAS. However, the broadcast capability on the

PON enables the AN (OLT) to send one copy on the PON as opposed to N copies of a multicast channel on the PON serving N premises being receivers. The PON multicast capability can be leveraged in the case of GPON and BPON as discussed in this document.

Fundamental to leveraging the broadcast capability on the PON for multicast delivery is the ability to assign a single encryption key for all PON frames carrying all multicast channels or a key per set of multicast channels that correspond to service packages, or none. It should be noted that the ONT can be a multi-Dwelling Unit (MDU) ONT with multiple Ethernet ports, each connected to a living unit. Thus, the ONT must not only be able to receive a multicast frame, but must also be able to forward that frame only to the Ethernet port with receivers for the corresponding channel.

In order to implement triple-play service delivery with necessary "quality-of-experience", including end-to-end bandwidth optimized multicast video delivery, there needs to be tight coordination between the NAS and the ANX. This interaction needs to be near real-time as services are requested via application or network level signaling by broadband subscribers. ANCP as defined in [ANCP-FRAMEWORK] for DSL based networks is very suitable to realize a control protocol (with transactional exchange capabilities), between PON enabled ANX and the NAS, and also between the components comprising the ANX i.e. between OLT and the ONT. Typical use cases for ANCP in PON environment include the following:

- Multicast
    - Optimized multicast delivery
    - Unified video resource control
    - NAS based provisioning of ANX
- Access topology discovery
- Remote connectivity check

5 Concept of Access Node Control Mechanism for PON based access

The high-level communication framework for an Access Node Control Mechanism is shown in Figure 3. The Access Node Control Mechanism defines a quasi real-time, general-purpose method for multiple network scenarios with an extensible communication scheme, addressing the different use cases that are described in the sections that follow. The access node control mechanism is also extended to run between OLT and ONT. The mechanism consists of control function, and reporting and/or enforcement function. Controller function is used to receive status information or admission requests from the reporting function. It is also used to trigger a certain behavior in the

network element where the reporting and/or enforcement function resides.

The reporting function is used to convey status information to the controller function that requires the information for executing local functions. The enforcement function can be contacted by the controller function to enforce a specific policy or trigger a local action. The messages shown in Figure 3 show the conceptual message flow. The actual use of these flows, and the times or frequencies when these messages are generated depend on the actual use cases, which are described in later sections.

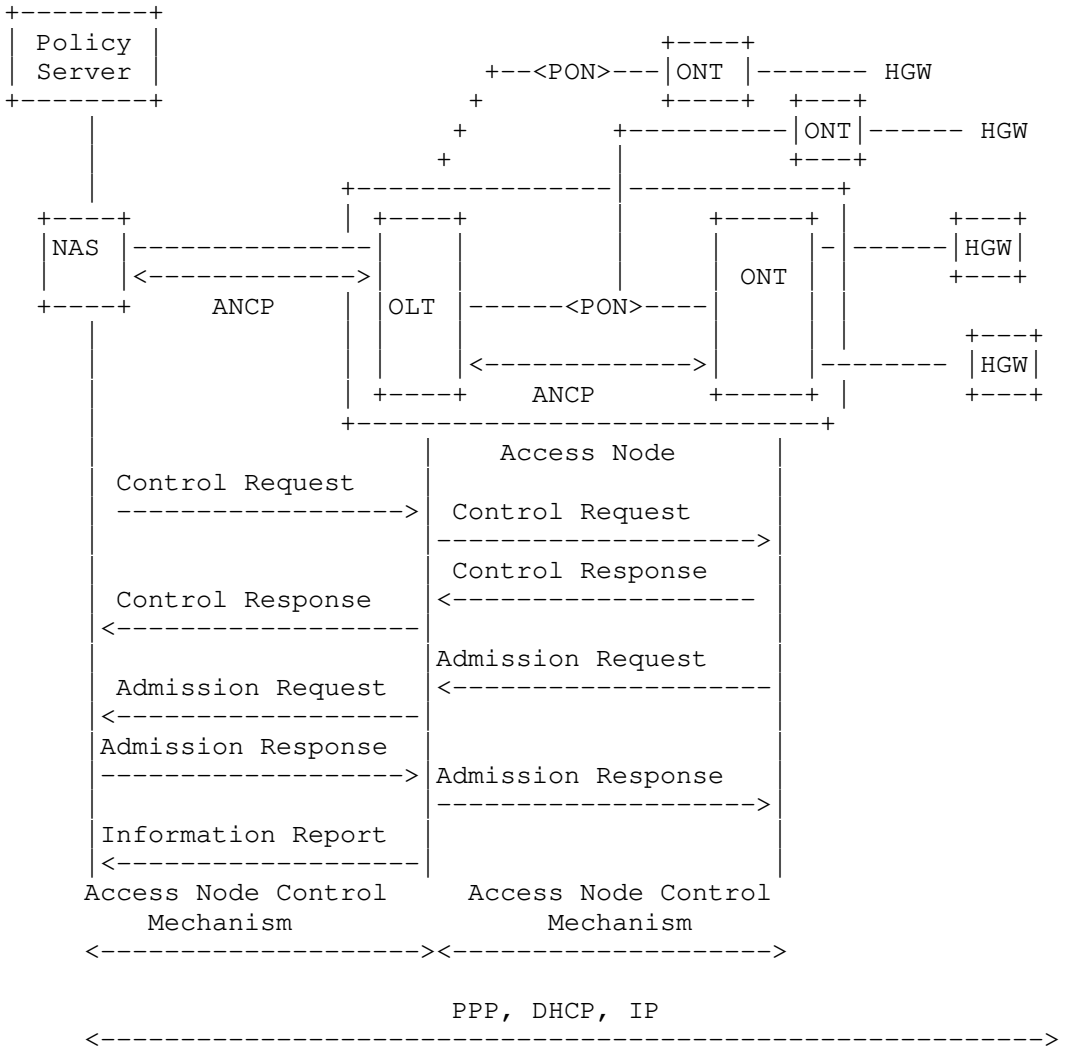

Figure 3. Conceptual Message Flow for Access Node Control Mechanism

6 Multicast

With the rise of supporting IPTV services in a resource-efficient way, multicast services are becoming increasingly important.

In order to gain bandwidth optimization with multicast, the replication of multicast content per access-loop needs to be distributed to the ANX. This can be done by ANX (OLT and ONT) becoming multicast aware by implementing an IGMP snooping and/or proxy function. The replication thus needs to be distributed between NAS, aggregation nodes, and ANX. In case of GPON, and in case of BPON with Ethernet uplink, this is very viable. By introducing IGMP processing on the ANX and aggregation nodes, the multicast replication process is now divided between the NAS, the aggregation node(s) and ANX. This is in contrast to the ATM-based model, where NAS is the single element responsible for all multicast control and replication. In order to ensure backward compatibility with the ATM-based model, the NAS, aggregation node and ANX need to behave as a single logical device. This logical device must have exactly the same functionality as the NAS in the ATM access/aggregation network. The Access Node Control Mechanism can be used to make sure that this logical/functional equivalence is achieved by exchanging the necessary information between the ANX and the NAS.

An alternative to multicast awareness in the ANX is for the subscriber to communicate the IGMP "join/leave" messages with the NAS, while the ANX is being transparent to these messages. In this scenario, the NAS can use ANCP to create replication state in the ANX for efficient multicast replication. The NAS sends a single copy of the multicast stream towards the ANX. The NAS can perform network-based conditional access and multicast admission control on multicast joins, and create replication state in the ANX if the request is admitted by the NAS.

The following sections describe various use cases related to multicast.

6.1 Multicast Conditional Access

In a Broadband FTTP access scenario, Service Providers may want to dynamically control, at the network level, access to some multicast flows on a per user basis. This may be used in order to differentiate among multiple Service Offers or to realize/reinforce

conditional access based on customer subscription. Note that, in some environments, application layer conditional access by means of Digital Rights Management (DRM) for instance may provide sufficient control, so that network-based Multicast conditional access may not be needed. However, network level access control may add to the service security by preventing the subscriber from receiving a non-subscribed channel. In addition, it enhances network security by preventing a multicast stream from being sent on a link or a PON based on a non-subscriber request.

Where network-based channel conditional access is desired, there are two approaches. It can be done on the NAS along with bandwidth based admission control. The NAS can control the replication state on the ANX based on the outcome of access and bandwidth based admission control. This is covered later in section 3.4. The other approach is to provision the necessary conditional access information on the ANX (ONT and/or OLT) so the ANX can perform the conditional access decisions autonomously. For these cases, the NAS can use ANCP to provision black and white lists as defined in [ANCP-FRAMEWORK], on the ANX so that the ANX can decide locally to honor a join or not. It should be noted that in the PON case, the ANX is composed of the ONT and OLT. Thus, this information can be programmed on the ONT and/or OLT. Programming this information on the ONT prevents illegitimate joins from propagating further into the network. A third approach, outside of the scope, may be to program the HGW with the access list.

A White list associated with an Access Port identifies the multicast channels that are allowed to be replicated to that port. A Black list associated with an Access Port identifies the multicast channels that are not allowed to be replicated to that port. It should be noted that the black list if not explicitly programmed is the complement of the white list and vice versa.

If the ONT performs IGMP snooping and it is programmed with a channel access list, the ONT will first check if the requested multicast channel is part of a White list or a Black list associated with the access port on which the IGMP join is received. If the channel is part of a White list, the ONT will pass the join request upstream towards the NAS. The ONT must not start replicating the associated multicast stream to the access port if such a stream is received until it gets confirmation that it can do so from the upstream node (NAS or OLT). Passing the channel access list is one of the admission control criteria whereas bandwidth-based admission control is another. If the channel is part of a Black list, the ONT

can autonomously discard the message because the channel is not authorized for that subscriber.

The ONT, in addition to forwarding the IGMP join, sends an ANCP admission request to the OLT identifying the channel to be joined and the premise. Premise identification to the OLT can be based on a Customer-Port-ID that maps to the access port on the ONT and known at the ONT and OLT. If the ONT has a white list and/or a black list per premise, the OLT need not have such a list. If the ONT does not have such a list, the OLT may be programmed with such a list for each premise. In this latter case, the OLT would perform the actions described earlier on the ONT. Once the outcome of admission control (conditional access and bandwidth based admission control) is determined by the OLT (either by interacting with the NAS or locally), it is informed to the ONT. OLT Bandwidth based admission control scenarios are defined in section 3.4.

The White List and Black List can contain entries allowing:

- An exact match for a (*,G) ASM group (e.g. <G=g.h.i.l>);
- An exact match for a (S,G) SSM channel (e.g. <S=s.t.u.v,G=g.h.i.l>);
- A mask-based range match for a (*,G) ASM group (e.g. <G=g.h.i.l/Mask>);
- A mask-based range match for a (S,G) SSM channel (e.g. <S=s.t.u.v,G=g.h.i.l/Mask>);

The use of a White list and Black list may be applicable, for instance, to regular IPTV services (i.e. Broadcast TV) offered by an Access Provider to broadband (e.g., FTTP) subscribers. For this application, the IPTV subscription is typically bound to a specific FTTP home, and the multicast channels that are part of the subscription are well-known beforehand. Furthermore, changes to the conditional access information are infrequent, since they are bound to the subscription. Hence the ANX can be provisioned with the conditional access information related to the IPTV service.

Instead of including the channel list(s) at the ONT, the OLT or NAS can be programmed with these access lists. Having these access lists on the ONT prevents forwarding of unauthorized joins to the OLT or NAS, reducing unnecessary control load on these network elements. Similarly, performing the access control at the OLT instead of the NAS, if not performed on the ONT, will reduce unnecessary control load on the NAS.

6.2 Multicast Admission Control

The successful delivery of Triple Play Broadband services is quickly becoming a big capacity planning challenge for most of the Service Providers nowadays. Solely increasing available bandwidth is not always practical, cost-economical and/or sufficient to satisfy end-user experience given not only the strict QoS requirements of unicast applications like VoIP and Video on Demand, but also the fast growth of multicast interactive applications such as "video conferencing", digital TV, and digital audio. These applications typically require low delay, low jitter, low packet loss and high bandwidth. These applications are also typically "non-elastic", which means that they operate at a fixed bandwidth, which cannot be dynamically adjusted to the currently available bandwidth.

An Admission Control (AC) mechanism covering admission of multicast traffic for the FTTP access is required, in order to avoid over-subscribing the available bandwidth and negatively impacting the end-user experience. Before honoring a user request to join a new multicast flow, the combination of ANX and NAS MUST ensure admission control is performed to validate that there is enough video bandwidth remaining on the PON, and on the uplink between the OLT and NAS to carry the new flow (in addition to all other existing multicast and unicast video traffic) and that there is enough video bandwidth for the subscriber to carry that flow. The solution needs to cope with multiple flows per premise and needs to allow bandwidth to be dynamically shared across multicast and unicast video traffic per subscriber, PON, and uplink (irrespective of whether unicast AC is performed by the NAS, or by some off-path Policy Server). It should be noted that the shared bandwidth between multicast and unicast video is under operator control. That is, in addition to the shared bandwidth, some video bandwidth could be dedicated to Video on Demand, while other video bandwidth could be dedicated for multicast. The focus in this document will be on multicast-allocated bandwidth including the shared unicast and multicast bandwidth. Thus, supporting admission control requires some form of synchronization between the entities performing multicast AC (e.g. the ANX and/or NAS), the entity performing unicast AC (e.g. the NAS or a Policy Server), and the entity actually enforcing the multicast replication (i.e., the NAS and the ANX). This synchronization can be achieved in a number of ways:

. - One approach is for the NAS to perform bandwidth based admission control on all multicast video traffic and unicast video traffic that requires using the shared bandwidth with multicast shr. Based on the outcome of admission control, NAS then controls the replication state on the ANX.

  The subscriber generates an IGMP join for the desired stream on its logical connection to the NAS. The NAS terminates the IGMP message, performs conditional access, and bandwidth based admission control on the IGMP request. The bandwidth admission control is performed against the following:

1. Available video bandwidth on the link to OLT

2. Available video bandwidth on the PON interface

3. Available video bandwidth on the last mile (access-port on the ONT/ONU).

The NAS can locally maintain and track video bandwidth it manages for all the three levels mentioned above. The NAS can maintain identifiers corresponding to the PON interface and the last mile (customer interface). It also maintains a channel map, associating every channel (or a group of channels sharing the same bandwidth requirement) with a data rate. For instance, in case of 1:1 VLAN representation of the premise, the outer tag (S-VLAN) could be inserted by the ANX to correspond to the PON interface on the OLT, and the inner-tag could be inserted by the ANX to correspond to the access-line towards the customer. Bandwidth tracking and maintenance for the PON interface and the last-mile could be done on these VLAN identifiers. In case if N:1 representation, the single VLAN inserted by ANX could correspond to the PON interface on the OLT. The access loop is represented via Customer-Port-ID received in "Agent Circuit Identifier" sub-option in DHCP messages.

The NAS can perform bandwidth accounting on received IGMP messages. The video bandwidth is also consumed by any unicast video being delivered to the CPE. NAS can perform video bandwidth accounting and control on both IGMP messages and on requests for unicast video streams when either all unicast admission control is done by the NAS or an external policy server makes a request to the NAS for using shared bandwidth with multicast as described later in the document.

This particular scenario assumes the NAS is aware of the bandwidth on the PON, and under all conditions can track the changes in available bandwidth on the PON. On receiving an IGMP Join message, NAS will perform bandwidth check on the subscriber bandwidth. If this passes, and the stream is already being forwarded on the PON by the OLT (which also means that it is already forwarded by the NAS to the OLT), NAS will admit the JOIN, update the available subscriber bandwidth, and transmit an ANCP message to the OLT and in turn to the ONT to start replication on the customer port. If the stream is not already being replicated to the PON by the OLT, the NAS will also check the available bandwidth on the PON, and if it is not already being replicated to the OLT it will check the bandwidth on the link towards the OLT. If this passes, the available PON bandwidth and the bandwidth on the link towards the OLT is updated. The NAS adds the OLT as a leaf to the multicast tree for that stream.

On receiving the message to start replication, the OLT will add the PON interface to its replication state if the stream is not already being forwarded on that PON. Also, the OLT will send an ANCP message to direct the ONT to add or update its replication state with the customer port for that channel. The interaction between

ANX and NAS is shown in Figures 4 and 5.
For unicast video streams, application level signaling from the CPE typically triggers an application server to request bandwidth based admission control from a policy server. The policy server can in turn interact with the NAS to request the bandwidth for the unicast video flow if it needs to use shared bandwidth with multicast. If the bandwidth is available, NAS will reserve the bandwidth, update the bandwidth pools for subscriber bandwidth, the PON bandwidth, and the bandwidth on the link towards the OLT, and send a response to the policy server, which is propagated back to the application server to start streaming. Otherwise, the request is rejected.

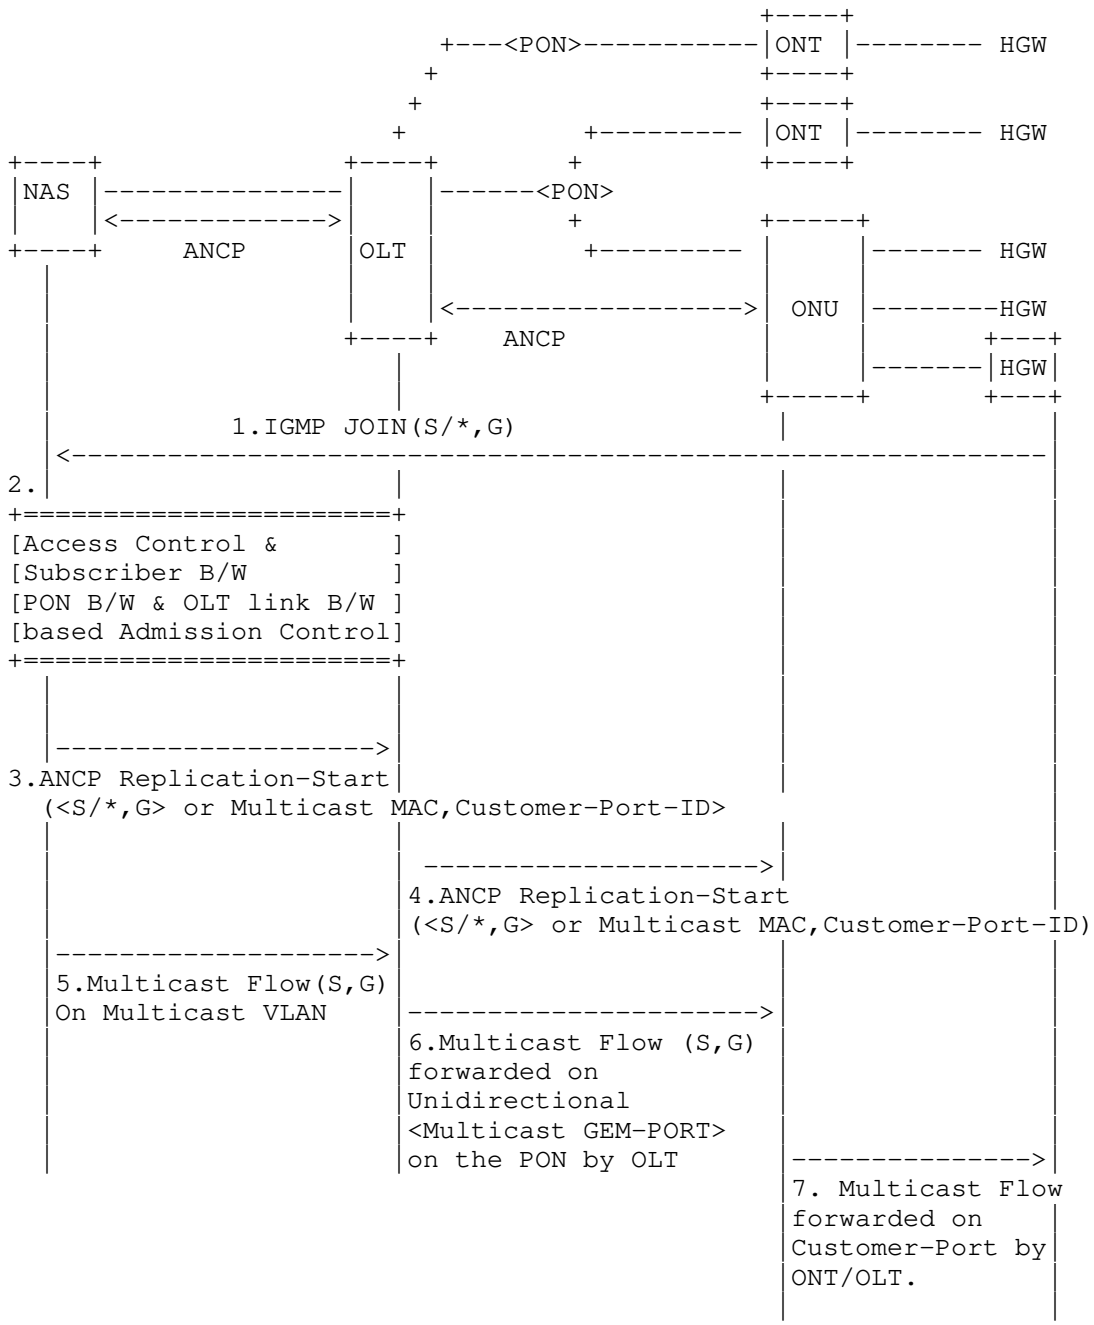

Figure 4. Interactions for NAS based Multicast Admission Control (no IGMP processing on ANX, and NAS maintains available video bandwidth for PON).

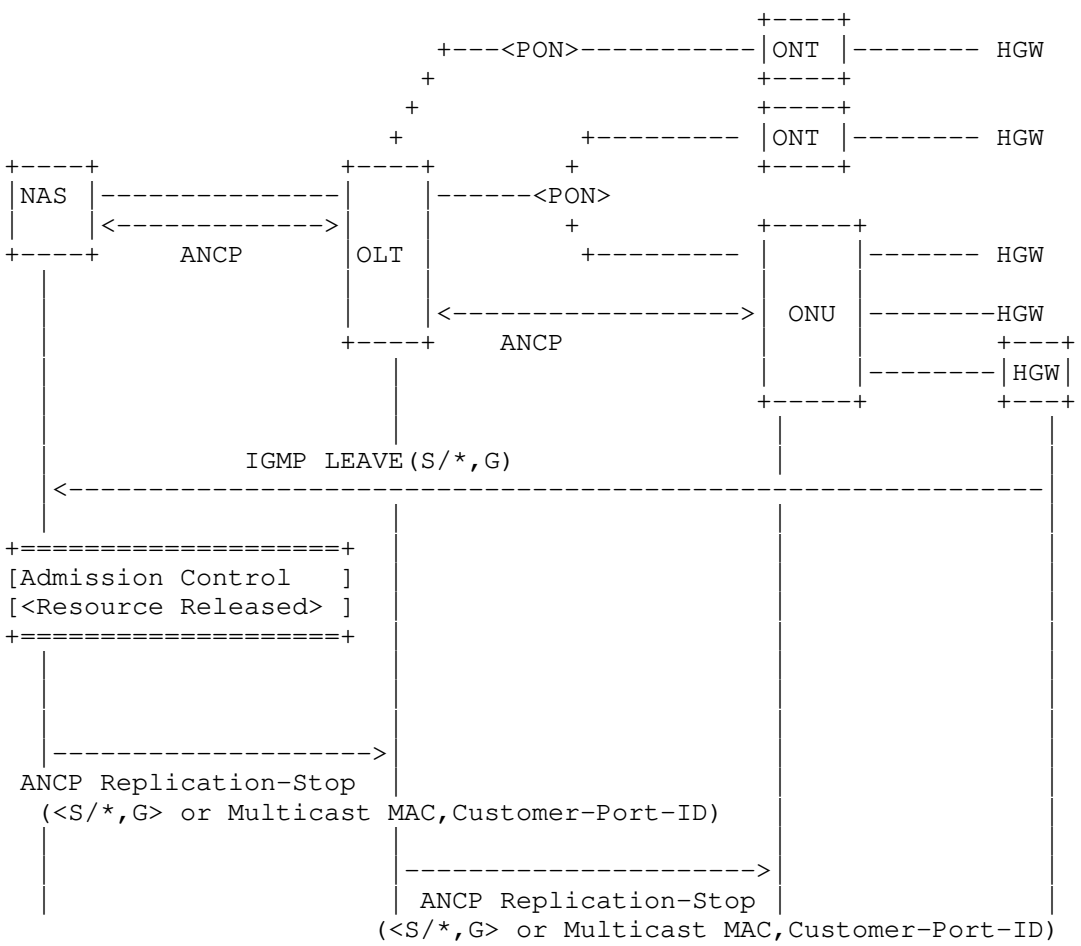

Figure 5. Interactions for NAS based Multicast Admission Control (no IGMP processing on ANX, and NAS maintains available video bandwidth for PON).

. An alternate approach is required if the NAS is not aware of the bandwidth on the PON. In this case the OLT does the PON bandwidth management, and requests NAS to perform bandwidth admission control on subscriber bandwidth and the bandwidth on the link to the OLT.

ANX operation:

- o ONT can snoop IGMP messages. If conditional access is configured and the channel is in the Black list (or it is not on the White list), ONT will drop the IGMP Join. If the channel passes the conditional access check, the ONT will forward the IGMP Join, and will send a bandwidth admission control request to the OLT. In case the multicast stream is already being received on the PON, the ONT does not forward the stream to the access port where IGMP is received, till it has received a positive admission control response from the OLT.

- o OLT can snoop IGMP messages. It also receives a bandwidth admission control request from the ONT for the requested channel. It can be programmed with a channel bandwidth map. If the multicast channel is already being streamed on the PON, or the channel bandwidth is less than the multicast available bandwidth on the PON, the OLT forwards the IGMP request to the NAS and keeps track of the subscriber (identified by customer-Port-ID) as a receiver. If the channel is not already being streamed on the PON, but the PON has sufficient bandwidth for that channel, the OLT reduces the PON multicast video bandwidth by the channel bandwidth and may optionally add the PON to the multicast tree without activation for that channel. This is biased towards a forward expectation that the request will be accepted at the NAS. The OLT forwards the IGMP join to the NAS. It also sends a bandwidth admission request to the NAS identifying the channel, and the premise for which the request is made. It sets a timer for the subscriber multicast entry within which it expects to receive a request from the NAS that relates to this request. If the PON available bandwidth is less than the bandwidth of the requested channel, the OLT sends an admission response (with a reject) to the ONT, and does not forward the IGMP join to the NAS.

NAS operation:

The NAS receives the IGMP join from the subscriber on the subscriber connection. When NAS receives the admission control request from ANX (also signifying the bandwidth on the PON is available), it performs admission control against the subscriber available multicast bandwidth. If this check passes, and the NAS is already transmitting that channel to the OLT, the request is accepted. If the check passes and the NAS is not transmitting the channel to the OLT yet, it performs admission control against the multicast video available bandwidth (this includes the dedicated

multicast bandwidth and the shared bandwidth between multicast and video on demand) on the link(s) to the OLT. If the check passes, the request is accepted, the available video bandwidth for the subscriber and downlink to the OLT are reduced by the channel bandwidth, and the NAS sends an ANCP admission control response (indicating accept) to the OLT, requesting the addition of the subscriber to the multicast tree for that channel. The OLT activates the corresponding multicast entry if not active and maintains state of the subscriber in the list of receivers for that channel. The OLT also sends an ANCP request to the ONT to enable reception of the multicast channel and forwarding to the subscriber access port. Otherwise, if the request is rejected, the NAS will send an admission reject to the OLT, which in turns removes the subscriber as a receiver for that channel (it it were added), and credits back the channel bandwidth to the PON video bandwidth if there is no other receiver on the PON for that channel. The interactions between ANX and NAS are show in Figures 6 and 7.

If the OLT does not receive a response from the NAS within a set timer, the OLT removes the subscriber from the potential list of receivers for the indicated channel. It also returns the allocated bandwidth to the PON available bandwidth if there are no other receivers. In this case, the NAS may send a response to the OLT with no matching entry as the entry has been deleted. The OLT must perform admission control against the PON available bandwidth and may accept the request and send an ANCP request to the ONT to activate the corresponding multicast entry as described earlier. If it does not accept the request, it will respond back to the NAS with a reject. The NAS shall credit back the channel bandwidth to the subscriber. It shall also stop sending the channel to the OLT if that subscriber was the last leaf on the multicast tree towards the OLT.

On processing an IGMP leave, the OLT will send an ANCP request to NAS to release resources. NAS will release the subscriber bandwidth. If this leave causes the stream to be no longer required by the OLT, the NAS will update its replication state and release the bandwidth on the NAS to OLT link.

If the subscriber makes a request for a unicast video stream (i.e., Video on Demand), it results in appropriate application level signaling, which typically results in an application server requesting a policy server for bandwidth-based admission control for the VoD stream. The policy server after authorizing the request, can send a request to the NAS for the required bandwidth if it needs to use bandwidth that is shared with multicast. This request may be based on a protocol outside of the scope of this document. The NAS checks if the available video bandwidth (accounting for both multicast and unicast) per subscriber and for the link to the OLT is sufficient for the request. If it is, it temporarily reserves the bandwidth and sends an ANCP admission request to the OLT for the subscriber, indicating the desired VoD bandwidth. If the OLT has sufficient bandwidth on the corresponding

PON, it reserves that bandwidth and returns an admission response to the NAS. If not, it returns a reject to the NAS. If the NAS receives an accept, it returns an accept to the policy server which in turn returns an accept to the application server, and the video stream is streamed to the subscriber. This interaction is shown in Figure 8. If the NAS does not accept the request from the policy server, it returns a reject. If the NAS receives a reject from the OLT, it returns the allocated bandwidth pool to the subscriber and the downlink to the OLT.

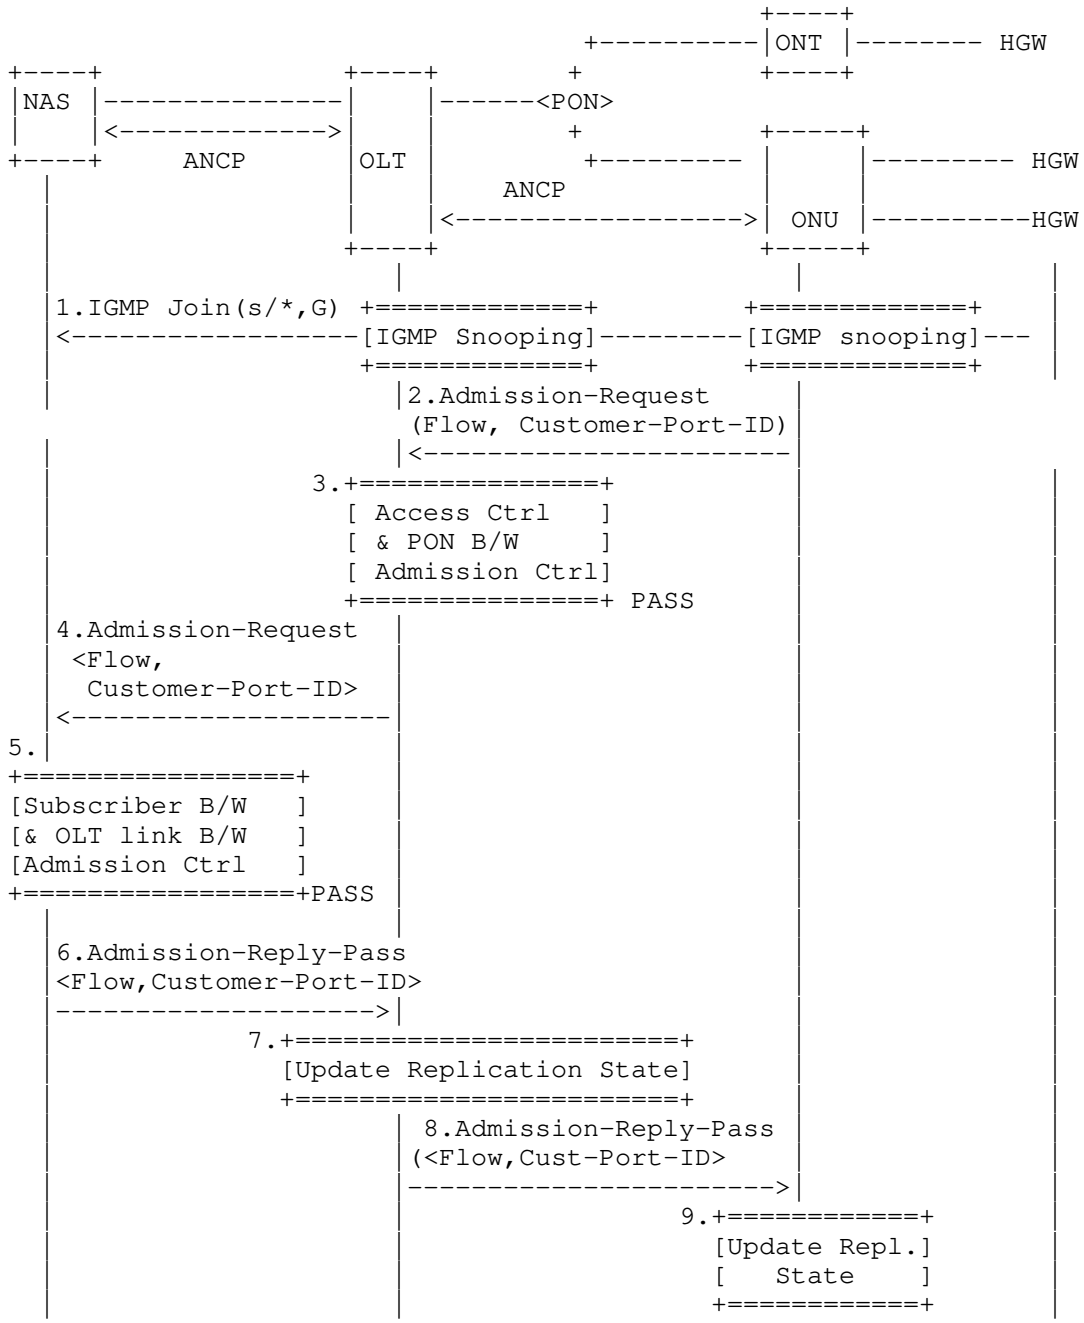

Figure 6. Interaction between NAS & ANX for Multicast B/W Admission Control

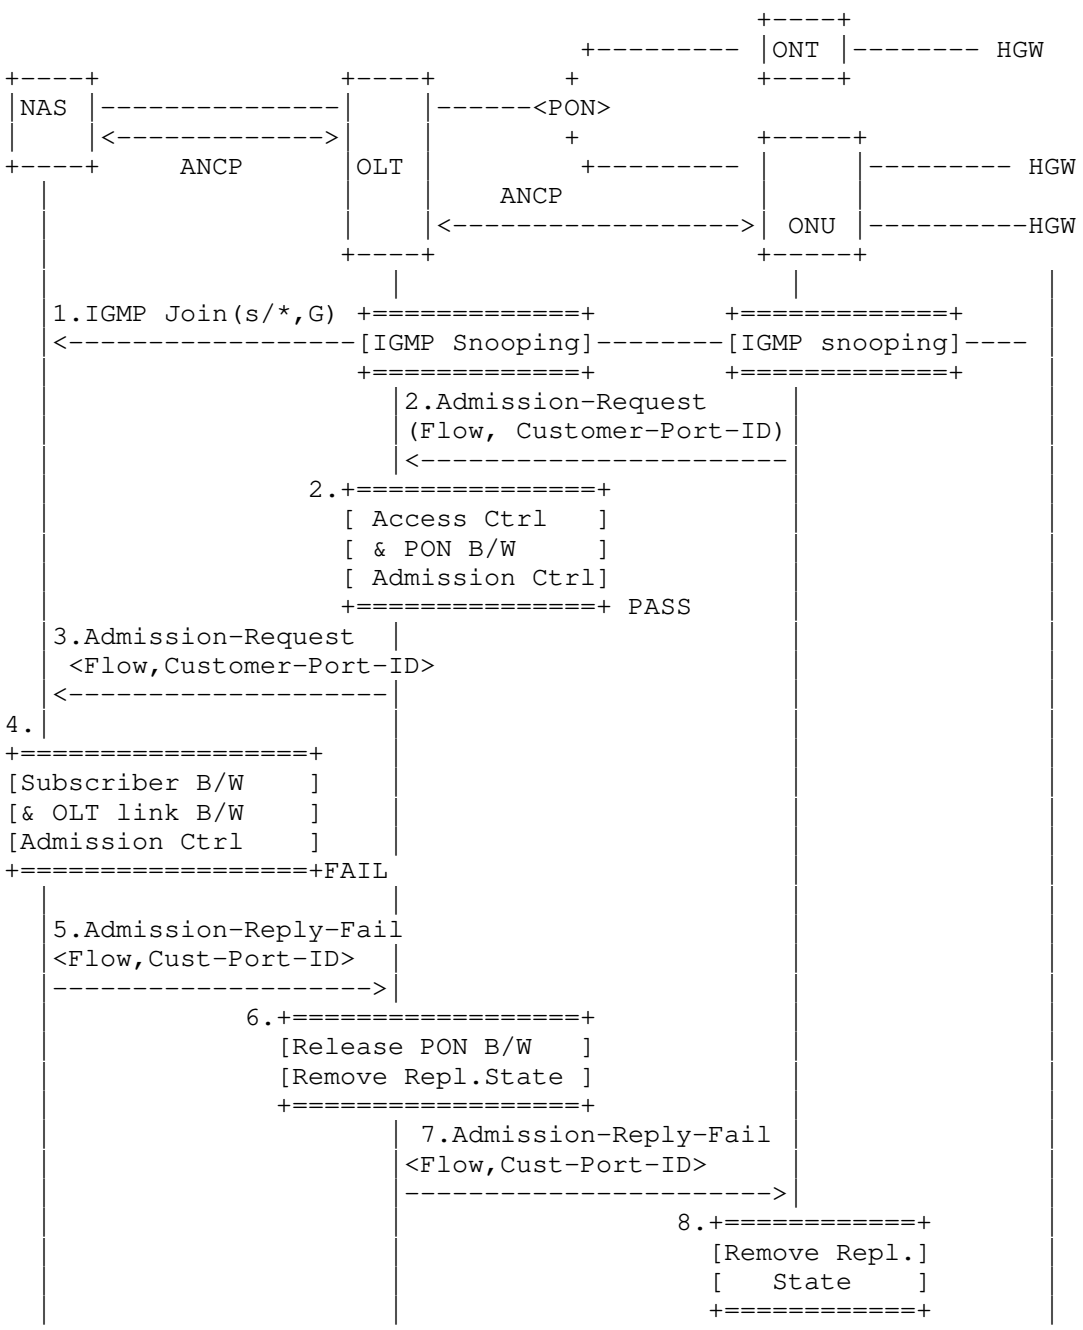

Figure 7. Interaction between NAS and ANX for Multicast B/W Admission Control

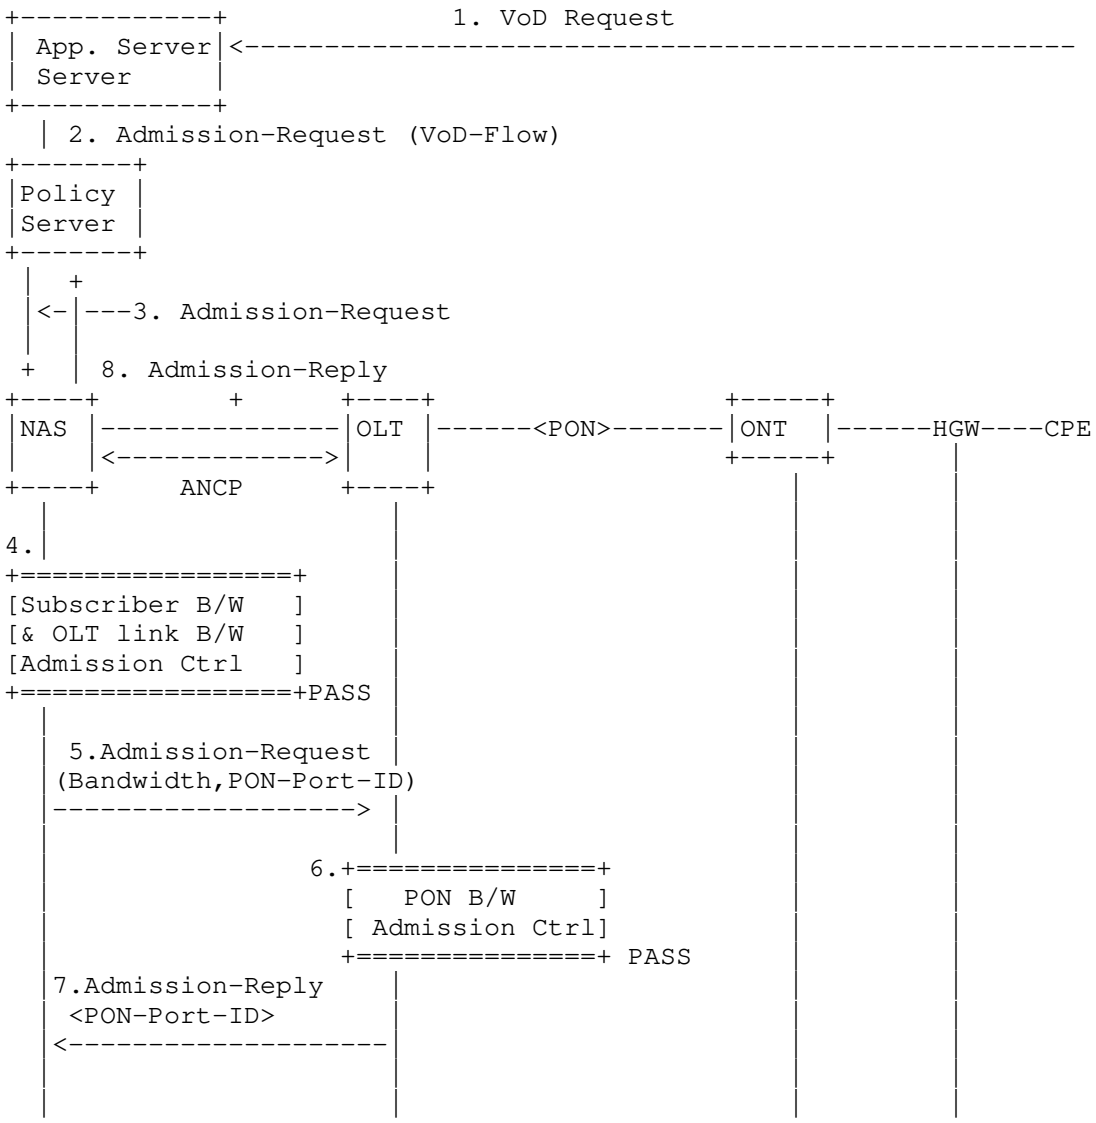

Figure 8. Interactions for VoD Bandwidth Admission Control

. A third possible approach is where the ANX is assumed to have a full knowledge to make an autonomous decision on admitting or rejecting a multicast and a unicast join. With respect to the interaction between ONT and OLT, the procedure is similar to the first approach (i.e. NAS controlled replication). However, when the OLT receives an IGMP request for a subscriber, it performs admission control against that subscriber multicast video bandwidth (dedicated and shared with Video on Demand), the PON and uplink to

the GWR. It should be noted in this case that if there are multiple NAS-OLT links, either the link on which the multicast stream must be sent is pre-determined, needs to be selected by the OLT based on downstream bandwidth from NAS to OLT and the selection is communicated to the NAS, or the OLT has to be ready to receive the stream on any link. If the check passes, the OLT updates the video available bandwidth per PON and subscriber. The OLT adds the subscriber to the list of receivers and the PON to the multicast tree, if it is not already on it. It also sends an ANCP request to the ONT to add the subscriber access port to that channel multicast tree, and sends an ANCP message to the NAS informing it of the subscriber and link available video bandwidth and the channel the subscriber joined. The NAS upon receiving the ANCP information message, updates the necessary information, including the OLT to the multicast tree if it is not already on it. It should be noted in this case that the ANCP message from the OLT to the NAS is being used to add the OLT to a multicast tree as opposed to an IGMP message. The IGMP message can also be sent by the OLT with the OLT acting as an IGMP proxy at the expense of added messages. In this option, the OLT acts as the network IGMP router for the subscriber.

For unicast video streams, the policy server receiving an admission request from an application server, as described before, may query the OLT for admission control as it has all information. If the OLT has sufficient bandwidth for the stream it reserves that bandwidth for the subscriber, PON and OLT uplink to the NAS and returns an accept to the policy server. It also updates the NAS via an ANCP message of the subscriber available video bandwidth. If the OLT rejects the policy server request, it will return a reject to the policy server.

It should be noted that if the policy server adjacency is with the NAS, the policy server may make the admission request to the NAS. The NAS then sends an ANCP admission request to the OLT on behalf of the policy server. The NAS returns an accept or reject to the policy server if it gets a reject or accept, respectively, from the OLT.

6.3 Multicast Accounting

It may be desirable to perform accurate per-user or per Access Loop time or volume based accounting. In case the ANX is performing the traffic replication process, it knows when replication of a multicast flow to a particular Access Port or user starts and stops. Multicast accounting can be addressed in two ways:

- ANX keeps track of when replication starts or stops, and reports this information to the NAS for further processing. In this case, ANCP can be used to send the information from the ANX to the NAS. This can be done with the Information Report message. The NAS can then generate the appropriate time and/or volume accounting information per Access Loop and per multicast flow, to be sent to the accounting system. The ANCP requirements to support this approach are specified in [ANCP-FRAMEWORK. If

the replication function is distributed between the OLT and ONT a query from the NAS will result in OLT generating a query to the ONT.

- o ANX keeps track of when replication starts or stops, and generates the time and/or volume based accounting information per Access Loop and per multicast flow, before sending it to a central accounting system for logging. Since ANX communicates with this accounting system directly, the approach does not require the use of ANCP. It is therefore beyond the scope of this document;

It may also be desirable for the NAS to have the capability to asynchronously query the ANX to obtain an instantaneous status report related to multicast flows currently replicated by the ANX. Such a reporting functionality could be useful for troubleshooting and monitoring purposes. If the replication function in the ANX is distributed between the OLT and the ONT, then for some of the information required by the NAS (such as the list of access-ports on which a flow is being forwarded or list of flows being forwarded on an access-port), a query to the OLT from the NAS will result in a query from OLT to ONT. The OLT responds back to the NAS when it receives the response from the ONT. Also, if the list of PONs on which replication is happening for a multicast channel or the list of channels being replicated on a PON is what is desired, the OLT can return this information.

7 Remote Connectivity Check

In an end-to-end Ethernet aggregation network, end-to-end Ethernet OAM as specified in IEEE 802.1ag and ITU-T Recommendation Y.1730/1731 can provide Access Loop connectivity testing and fault isolation. However, most HGWs do not yet support these standard Ethernet OAM procedures. Also, in a mixed Ethernet and ATM access network (e.g. Ethernet based aggregation upstream from the OLT, and BPON downstream), interworking functions for end-to-end OAM are not yet standardized and widely available. Until such mechanisms become standardized and widely available, Access Node Control mechanism between NAS and ANX can be used to provide a simple mechanism to test connectivity of an access-loop from the NAS.

Triggered by a local management interface, the NAS can use the Access Node Control Mechanism (Control Request Message) to initiate an Access Loop test between Access Node and HGW. On reception of the ANCP message, the OLT can trigger native OAM procedures defined for BPON in [G.983.1] and for GPON in [G.984.1]. The Access Node can send the result of the test to the NAS via a Control Response message.

8 Access Topology Discovery

In order to avoid congestion in the network, and manage and utilize the network resources better, and ensure subscriber fairness, NAS performs hierarchical shaping and scheduling of the traffic by modeling different congestion points in the network (such as the last-mile, Access Node uplink, and the access facing port).

Such mechanisms require that the NAS gains knowledge about the topology of the access network, the various links being used and their respective rates. Some of the information required is somewhat dynamic in nature (e.g. DSL line rate in case the last mile is xDSL based e.g. in case of "PON fed DSLAMs" for FTTC/FTTN scenarios), hence cannot come from a provisioning and/or inventory management OSS system. Some of the information varies less frequently (e.g. capacity of the OLT uplink), but nevertheless needs to be kept strictly in sync between the actual capacity of the uplink and the image the NAS has of it.

OSS systems are rarely able to enforce in a reliable and scalable manner the consistency of such data, notably across organizational boundaries under certain deployment scenarios. The Access Topology Discovery function allows the NAS to perform these advanced functions without having to depend on an error-prone and possibly complex integration with an OSS system.

The rate of the access-loop can be communicated via ANCP (Information Report Message) from the ONT to the OLT, and from OLT to the NAS. Additionally, during the time the DSL NT is active, data rate changes can occur due to environmental conditions (the DSL Access Loop can get "out of sync" and can retrain to a lower value, or the DSL Access Loop could use Seamless Rate Adaptation making the actual data rate fluctuate while the line is active). In this case, ANX sends an additional Information Report to the NAS each time the Access Loop attributes change above a threshold value.

9 Security Considerations

[ANCP-SECURITY] lists the ANCP related security threats that could be encountered on the Access Node and the NAS. It develops a threat model for ANCP security, and lists the security functions that are required at the ANCP level.

With Multicast handling as described in this document, ANCP protocol activity between the ANX and the NAS is triggered by join/leave requests coming from the end-user equipment. This could potentially be used for denial of service attack against the ANX and/or the NAS.

To mitigate this risk, the NAS and ANX MAY implement control plane protection mechanisms such as limiting the number of multicast flows a given user can simultaneously join, or limiting the maximum rate of join/leave from a given user.

Protection against invalid or unsubscribed flows can be deployed via provisioning black lists as close to the subscriber as possible (e.g. in the ONT).

10 Differences in ANCP applicability between DSL and PON

As it currently stands, both ANCP framework [ANCP-FRAMEWORK] and protocol [ANCP-PROTOCOL] are defined in context of DSL access. Due to inherent differences between PON and DSL access technologies, ANCP needs a few extensions for supporting the use-cases outlined in this document for PON based access. These specific differences and extensions are outlined below.

o In PON, the access-node functionality is split between OLT and ONT. Therefore, ANCP interaction between NAS and AN translates to transactions between NAS and OLT and between OLT and ONT. The processing of ANCP messages (e.g. for multicast replication control) on the OLT can trigger generation of ANCP messages from OLT to ONT. Similarly, ANCP messages from ONT to the OLT can trigger ANCP exchange between the ONT and the NAS (e.g. admission-request messages).This is illustrated in the generic message flow in Figure 3 of section 5. In case of DSL, the ANCP exchange is contained between two network elements (NAS and the DSLAM).

o The PON connection to the ONT is a shared medium between multiple ONTs on the same PON. The local-loop in case of DSL is point-to-point. In case of DSL access network, the access facing port on the NAS (i.e. port to the network between NAS and the DSLAM), and the access-facing ports on the DSLAM (i.e. customer's local-loop) are the two bandwidth constraint points that need to be considered for performing bandwidth based admission control for multicast video and VOD delivered to the customer. In case of PON access, in addition to the bandwidth constraint on the NAS to OLT facing ports, and the subscriber allocated bandwidth for video services, the bandwidth available on the PON for video is an additional constraint that needs to be considered for bandwidth based admission control. If the bandwidth control is centralized in NAS (as described in option 1 of section 6.2), then the NAS needs to support additional logic to consider available PON bandwidth before admitting a multicast request or a VOD request by the user. Accordingly, ANCP needs to identify the customer access port and the PON on which the customer ONT is. If the PON bandwidth control is performed on the OLT (as defined in second option in section 6.2), then additional ANCP request and response messages are required for NAS to query the OLT to determine available PON bandwidth when a request to admit a VOD flow is received on the NAS (as shown in figure 8 in section 6.2) or for the OLT to inform the NAS what stream bandwidth is sent to the subscriber for the NAS to take appropriate action (e.g., bandwidth adjustment for various types of traffic).

o In PON, the multicast replication can potentially be performed on three different network elements: (1) on the NAS (2) on the OLT for

replication to multiple PON ports and (3) on the ONT/ONU for replication to multiple customer ports. In case of DSL, the replication can potentially be performed on NAS and/or the DSLAM. Section 6.2 defines options for multicast replication in case of PON. In the first option, the multicast replication is done on the AN, but is controlled from NAS via ANCP (based on the reception of per-customer IGMP messages on the NAS). In this option, the NAS needs to supply to the OLT the set of PON-customer-IDs (as defined in section 2.1) to which the multicast stream needs to be replicated. The PON-customer-ID identifies the OLT and the PON ports on the OLT as well as the ONT and the access-ports on the ONT where the multicast stream needs to be replicated. Upon receiving the request to update its multicast replication state, the OLT MUST update its replication state with the indicated PON ports, but MAY also need to interact with the ONT via ANCP to update the multicast replication state on the ONT with the set of access-ports (as indicated by the NAS).In case of DSL, the DSLAM only needs to update its own replication state based on the set of access-ports indicated by the NAS.

o For reporting purposes, ANCP must enable the NAS to query the OLT for channels replicated on a PON or a list of PONs and to specific access ports. The latter should trigger the OLT to query the ONT for a list of channels being replicated on all access ports or on specific access ports to the premise. In DSL case, it is sufficient to query the DSLAM for a list of channels being replicated on an access port or a list of access ports.

11 ANCP versus OMCI between the OLT and ONT

ONT Management and Control Interface (OMCI) [OMCI] is specified for in-band ONT management via the OLT. This includes configuring parameters on the ONT. Such configuration can include adding an access port on the ONT to a multicast tree and the ONT to a multicast tree. Thus, OMCI can be a potential replacement for ANCP between the OLT and ONT, albeit it may not be suitable protocol for dynamic transactions as required for the multicast application.

If OMCI is selected to be enabled between the OLT and ONT to carry the same information elements that would be carried over ANCP, the OLT must perform the necessary translation between ANCP and OMCI for replication control messages received via ANCP. OMCI is an already available control channel, while ANCP requires a TCP/IP stack on the ONT that can be used by an ANCP client and accordingly it requires that the ONT be IP addressable for ANCP. Most ONTs today have a TCP/IP stack used by certain applications (e.g., VoIP, IGMP snooping). ANCP may use the same IP address that is often assigned to SIP or depending on the implementation may require a different address. Sharing the same IP address between SIP and ANCP may have other network implications on traffic routing. Using a separate IP address for the purpose of ONT management or ANCP specifically may often be required when supporting ANCP. These considerations may favor OMCI in certain environments. However, OMCI will not allow some of the transactions required in approach 2, where the ONT sends unsolicited requests to the OLT rather than being queried or

configured by OLT requests.

12 IANA Considerations

This document does not require actions by IANA.

13 Acknowledgements

14 References

14.1 Normative References

[RFC2516] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, February 1999.

[RFC2684] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 2684, September 1999.

14.2 Informative References

[RFC2881] Mitton, D. and M. Beadles, "Network Access Server Requirements Next Generation (NASREQNG) NAS Model", RFC 2881, Jul 2000.

[ANCP-FRAMEWORK] Ooghe, S., et al., "Framework and Requirements for Access Node Control Mechanism in Broadband Networks", RFC 5851, May 2010.

[G.983.1] ITU-T recommendation G.983.1, Broadband optical access systems based on Passive Optical Networks (PON).

[G.984.1] ITU-T recommendation G.984.1 Gigabit-capable Passive Optical Networks (G-PON): General characteristics

[TR-101] Cohen, A. and E. Shrum, "Migration to Ethernet-Based DSL Aggregation", DSL Forum TR-101, May 2006.

[ANCP-SECURITY] Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", RFC 5713, January 2010.

[OMCI] ITU-T recommendation G.984.4 GPON ONT Management and Control Interface (OMCI) Specifications.

[ANCP-PROTOCOL] Wadhwa, S et al, "Protocol for Access Node Control Mechanism in Broadband Networks", draft-ietf-ancp-protocol-12.txt, August 2010, work in progress.

Author's Addresses

Nabil Bitar
Verizon
117 West Street
Waltham, MA 02451

Email: nabil.n.bitar@verizon.com

Sanjay Wadhwa
Juniper Networks
10 Technology Park Drive
Westford, MA 01886

Email: swadhwa@juniper.net