ANCP
Internet-Draft
Updates: 6320 (if approved)
Intended status: Standards Track
Expires: August 29, 2014

F. Le Faucheur
Cisco
R. Maglione
Cisco Systems
T. Taylor
Huawei

February 25, 2014

Multicast Control Extensions for ANCP
draft-ietf-ancp-mc-extensions-16.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 and admission control with white and black lists;
- conditional access and admission control with grey lists;
- bandwidth delegation;
- committed bandwidth reporting.

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

This document updates RFC 6320 by assigning capability type 3 to a capability specified in this document and by changing the starting point for IANA allocation of result codes determined by IETF Consensus from 0x100 to 0x64.

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. [RFC6320] specifies a protocol for Access Node Control in broadband networks in line with this framework.

[RFC6320] 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 and admission control with white and black lists;
- conditional access and admission control 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 Section 3.4.1 of [RFC5851]. Section 3.4.2.2 mentions a way in which conditional access can be combined with admission control to allow best effort multicast flows. Section 3.4.2.3 points out the necessary conditions for using both conditional access and admission control.

In the case of "conditional access and admission control 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 combine conditional access with admission control of white-listed flows by appropriate provisioning.

Conditional access and admission control with grey lists is similar to conditional access and admission control 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.

This document updates RFC 6320 by assigning capability type 3 to the NAS-initiated multicast replication capability and by changing the starting point for IANA allocation of result codes determined by IETF Consensus from 0x100 to 0x64.

1.1. A Note On Scope

The requirements in [RFC5851] were formulated with the IPTV application in mind. Two basic assumptions underlie the use case descriptions:

- that the Home Gateway operates in bridged mode, and
- that multicast signalling uses IGMP ([RFC2236] or [RFC3376]) or MLD [RFC3810] rather than PIM [RFC4601].

Without the first assumption the AN may lose sight of individual subscriber devices making requests for multicast service. This has a very minor effect on the capabilities described below, but prevents the application of per-device policies at the NAS. Changing the second assumption would require that, in applications where the AN is responsible for snooping IGMP and MLD, it now also monitor for PIM signalling. The capabilities described in the present document do not depend explicitly on what type of multicast signalling is used, but the multiple phases of PIM setup could add complexity to their implementation.

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].

This document reuses the terms "connection admission control" ("CAC" or simply "admission control") and "conditional access" as they are used in [RFC5851].

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 referred to [RFC5851] for a more exhaustive description of the ANCP multicast use cases. Detailed example message flows can also be found in Appendix A.

In the diagrams below, participation of the Home Gateway is optional, depending on whether it is operating in bridged or routed mode. Note that devices behind the Home Gateway may require the Home Gateway to operate in routed mode to ensure that they can obtain access to non-IPTV multicast services.

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 performs conditional access and 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.

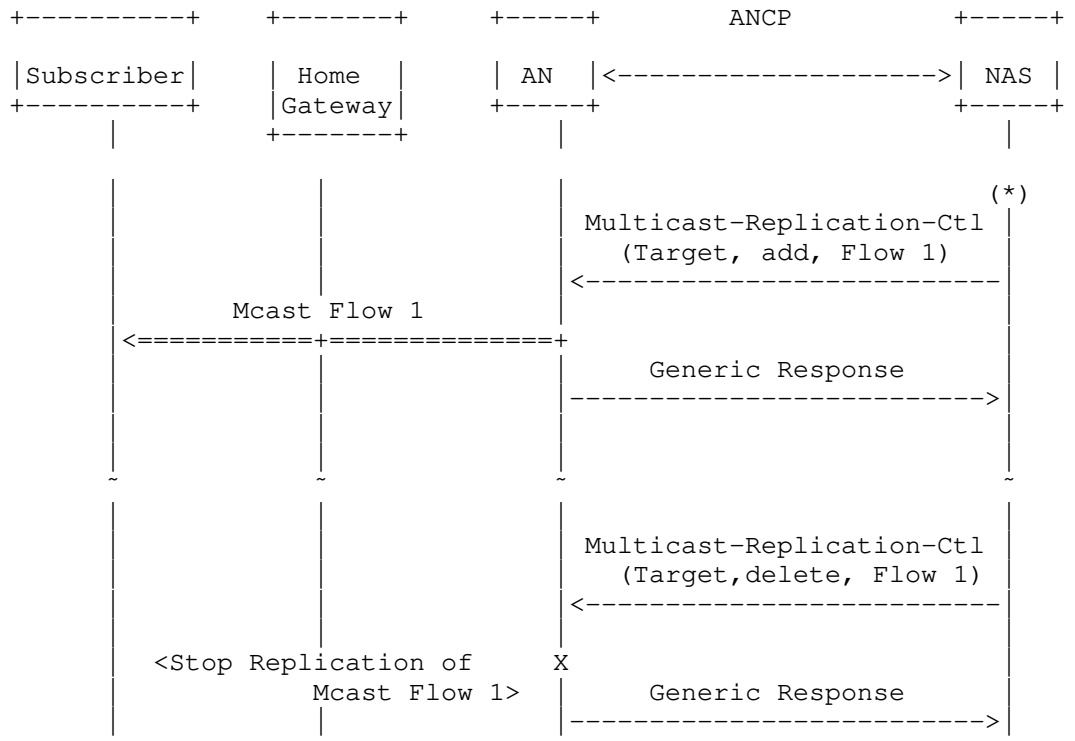

(*) The NAS may optionally seek direction from an external Authorization/Policy Server before admitting the flow.

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 end user for a multicast flow (in the grey list), the AN uses ANCP to request a 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 a conditional access and admission control check. In the case of a 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.

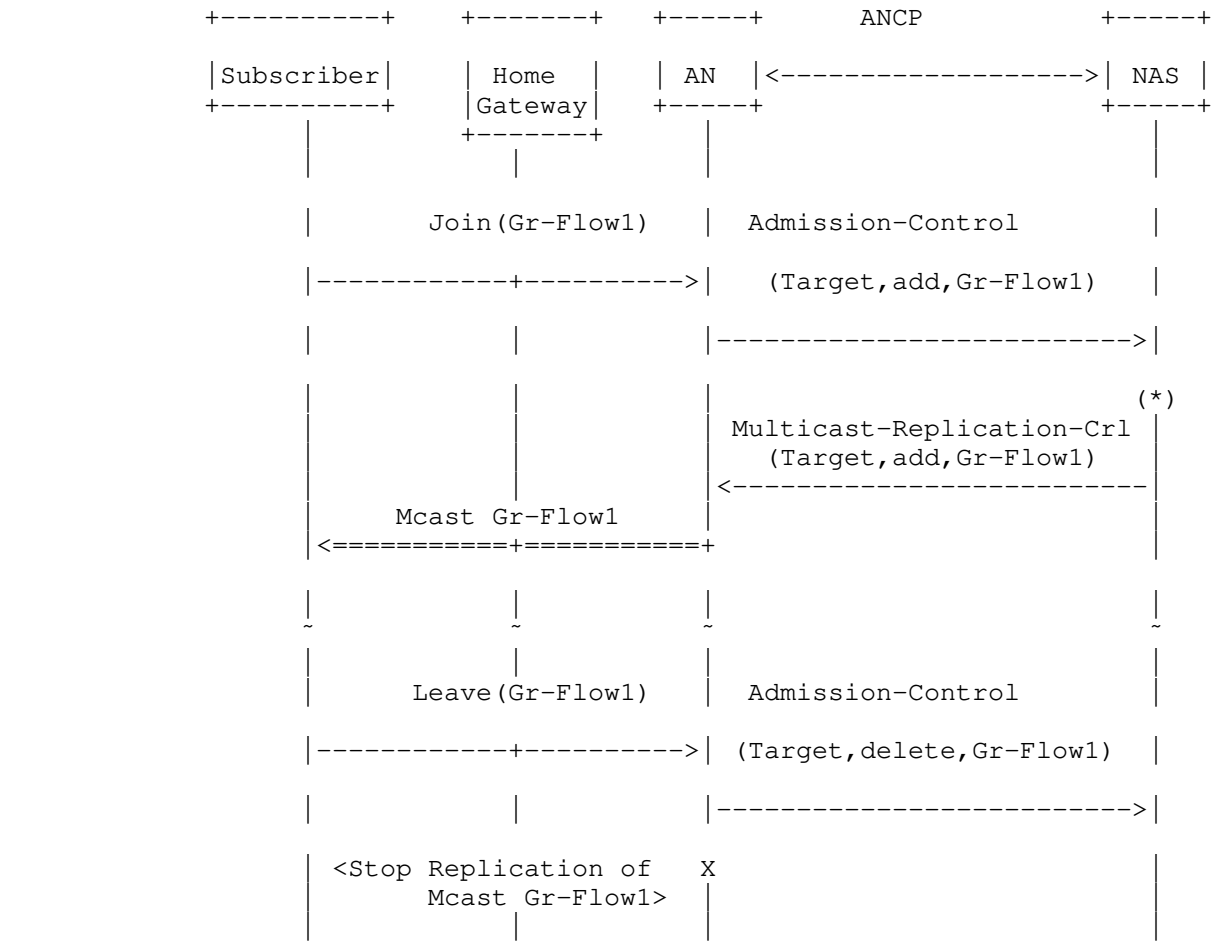

Gr-Flow1: a multicast flow matching the grey list for that port

(*) 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.

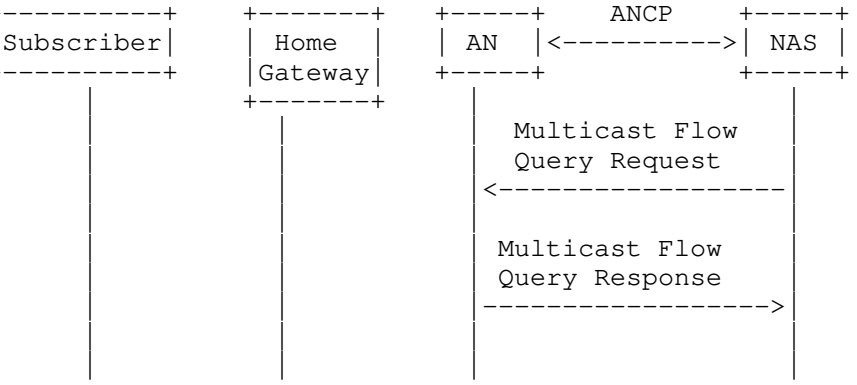

Figure 3: Multicast Flow Reporting

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.

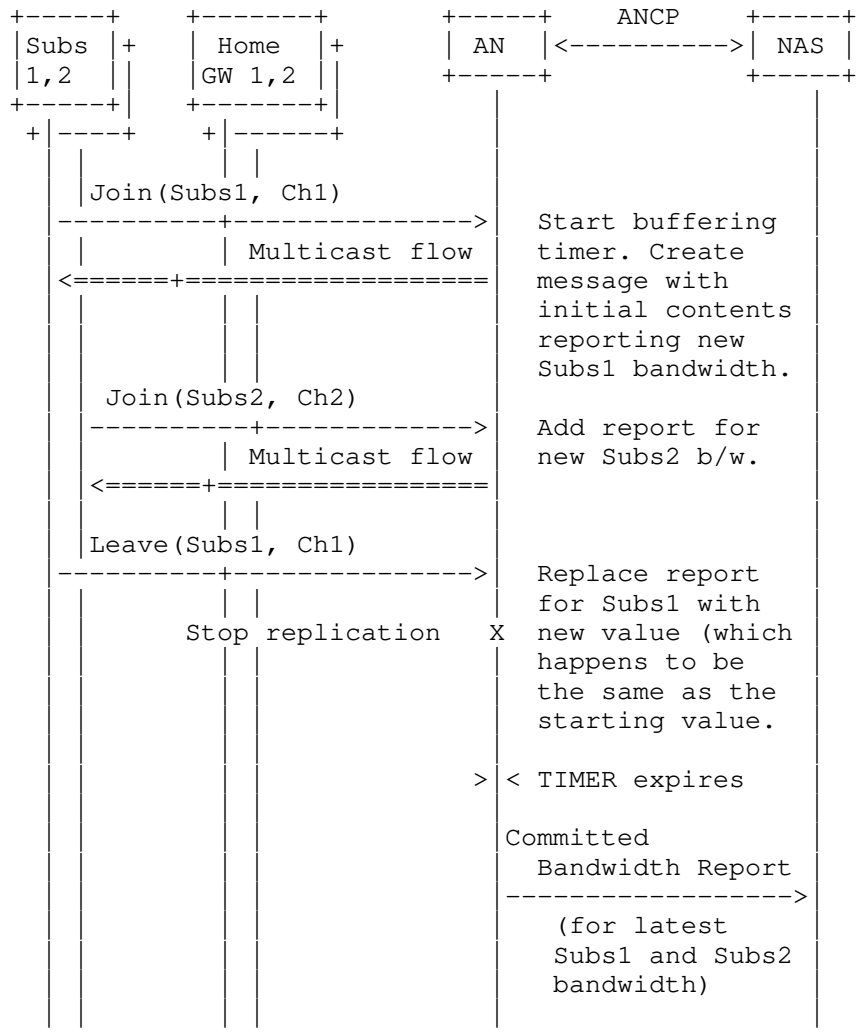

Figure 4: Message Flow For Committed Bandwidth Reporting

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.

Unless stated otherwise, receivers MUST ignore message contents that are not supported by the set of capabilities negotiated between the NAS and the Access Node.

4.1. Provisioning Message

Section 4.1 of [RFC6320] 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 5.13. 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 direct 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 MUST 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 and admission control 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 and admission control 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. In keeping with the general rule stated in

Section 4, 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 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 admission control on the indicated class(es) of flow.

The buffering time specified in an instance of the Report-Buffering-Time TLV will not be applied until the current accumulation process of Committed Bandwidth Report messages finishes.

As indicated in [RFC6320], 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 [RFC6320], 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 [RFC6320], the NAS may send DSL line configuration information to the AN ("ANCP based DSL Line Configuration" use case) using ANCP Port Management messages. See Section 7.3 of [RFC6320] for the format of the Port Management message in that usage.

This document specifies that the Port Management message MAY be used to convey 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 line specified by the extension block, and in the case of white and black lists, delegates conditional access to the AN for the specified access line and channels.

- 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 line.

When the Port Management message is used for this purpose:

- the Function field in the Port Management message MUST be set to 8, "Configure Connection Service Data".

- the message MUST include TLV(s) to identify the access line concerned. If the access line is a DSL loop, the line-identifying TLV(s) MUST be as specified in Section 5.1.2 of [RFC6320]. For non-DSL access lines, the appropriate alternative line-identifying TLV(s) MUST be present. Line configuration data other than the two TLVs listed in the previous paragraph MAY be present.

4.2.1. Sender Behaviour

The NAS sends the Port Management message at startup time to initialize parameters associated with the access line 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 line. This association replaces any previous association. That is, a given access line 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, unless such action is required by local policy. 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 Result 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 Result Code field in the header indicates 0x55, "Invalid TLV contents", 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 Section 4.3 of [RFC6320]. It MUST appear once and only once. It is encoded as specified in [RFC6320] or extensions and identifies the AN port subject to the request for admission or release.

- Command TLV: The Command TLV is defined in Section 4.4 of [RFC6320]. It MUST be present. It MAY appear multiple times.

As [RFC6320] 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.

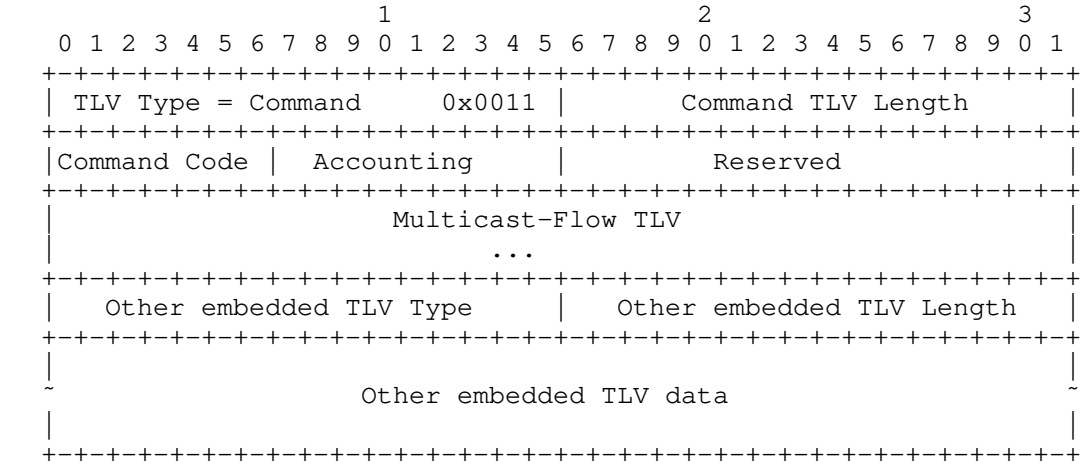

Figure 5: Contents of the Command TLV in the Multicast Replication Control Message

Command Code:

Command directive:

1 "Add";

2 "Delete";

3 "Delete All";

4 "Admission Control Reject";

5 "Conditional Access Reject";

6 "Admission Control and Conditional Access Reject".

Directives 4 through 6 are used as described in Section 4.4.2.

Accounting:
Meaningful only when the Command Code is "Add" (1). In that case, 0 indicates flow accounting is disabled, 1 indicates that octet accounting for the flow is requested. The sender MUST set the Accounting field to 0 and the receiver MUST ignore the Accounting field for other Command Code values.

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

Multicast-Flow TLV:
An instance of the Multicast-Flow TLV (Section 5.12) specifying the flow to be added or deleted. The Multicast-Flow TLV is omitted if the Command Code has value "Delete All" (3).

Other embedded TLV:
No other embedded TLVs are currently specified within the Multicast Replication Control message/Command TLV. However, see the description of the Multicast Admission Control message (Section 4.4). 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 Source-Specific Multicast (SSM) flows 2001:DB8::1, FF34::2, to 2001:DB8::2, FF34::3 on the Target identified by the "Access Loop Circuit ID":

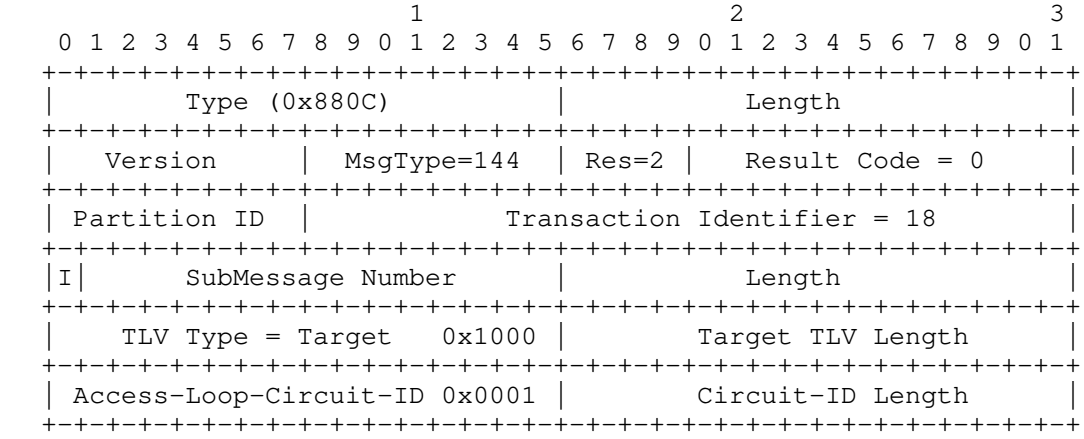

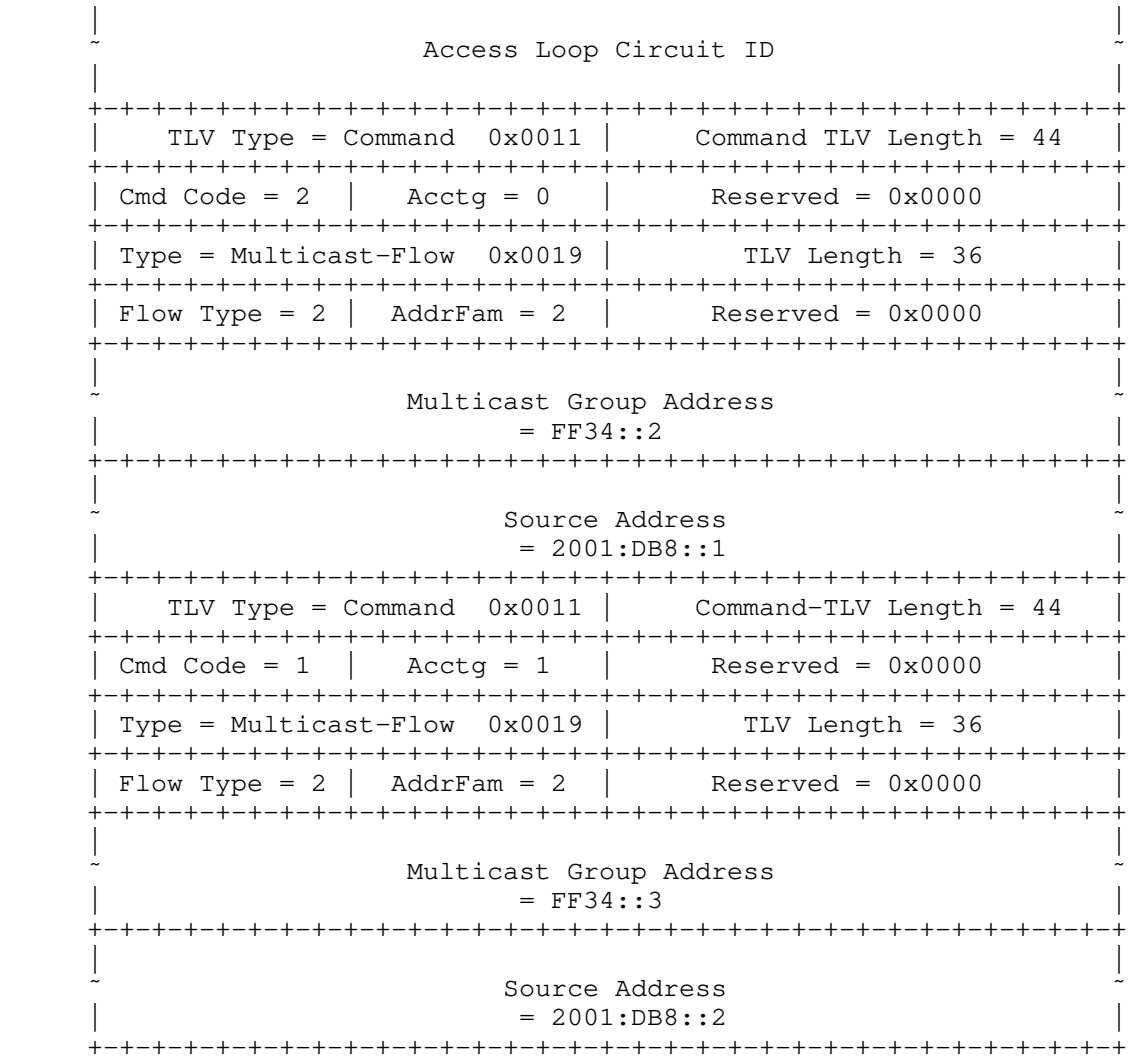

4.3.1. Sender Behaviour

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 Result Code field with the value 0 and the ANCP Transaction Identifier field with a unique value, as described in Section 3.6.1.6 of [RFC6320].

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, update the state of the accounting feature to what is set in the final command received, but all other features are unaffected by the second 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 1 for the first command.

If a Command TLV adds one or more flows and the AN is performing admission control for Multicast Replication Control messages, then the AN MUST perform admission control before replicating the flows. If the admission control check fails, the AN MUST treat the failure as an error as described below. The appropriate Result Code value for the response is 0x13 "Out of 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 Result Code field is set to 0, 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. Similarly, if a given command specifies multiple Single-Source Multicast (SSM) flows and a error occurs, processing MUST be interrupted at that point and the remainder of the Command TLV 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 Result Code values beyond those specified in [RFC6320], which MAY be used in a Generic Response sent in reply to a Multicast Replication Control message:

0x64 Command error.

Where detected: ANCP agent at the AN.

Further description: an invalid command code has been received.

Required additional information in the message: see below.

Target: ANCP agent at the NAS.

Action RECOMMENDED for the receiving ANCP agent: Report the error to the control application with an indication of the erroneous information associated with the invalid TLV(s).

0x65 Invalid flow address.

Where detected: ANCP agent at the AN.

Further description: either inconsistent flow address information has been provided or the address family is unsupported.

Required additional information in the message: see below.

Target: ANCP agent at the NAS.

Action RECOMMENDED for the receiving ANCP agent: Report the error to the control application with an indication of the erroneous information associated with the invalid TLV(s).

0x66 Multicast flow does not exist.

Where detected: control application at the AN.

Further description: the NAS has attempted to delete a flow that is not active on the given access line.

Required additional information in the message: see below.

Target: control application at the NAS.

Action RECOMMENDED for the receiving ANCP agent: report the error to the control application with an indication of the erroneous information associated with the invalid TLV(s).

A Generic Response message responding to the Multicast Replication Control message and containing one of the above Result 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 Result Code value in the message header. For example, the Result Code value could be 0x65 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 [RFC6320]. It MUST appear once and only once in the Multicast Admission Control message. It is encoded as specified in [RFC6320] or extensions and identifies the AN port subject to the request for admission or release.

- Command TLV: The Command TLV is defined in [RFC6320]. 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.

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 cannot 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.

- Finally and preferably, a Request-Source-Device-Id 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 3.6.1.6 of [RFC6320].

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 "Add" (1) when the message conveys a Join , to "Delete" (2) when the message conveys a Leave and to "Delete All" (3) 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 3, the Multicast-Flow TLV is 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.
- As a third alternative, the Request-Source-Device-Id embedded TLV MAY be included by the AN to convey a local identifier 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 inclusion of Request-Source-IP or Request-Source-MAC in the Multicast Admission Control message is typically done to allow the application of policies applicable to specific devices within the customer's network. However, transmission of either of these fields beyond the AN introduces potential privacy issues. Instead of transmitting either of these identifiers, it is RECOMMENDED that the AN map the required identifier to a local value known to the AN and AAA but not to the NAS, as discussed in Section 8. The local identifier is transmitted using the Request-Source-Device-Id TLV.

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 "Delete" (2) or "Delete All" (3) 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 "Add" (1) 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 with a unique value, as described in Section 3.6.1.6 of [RFC6320];
  - 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 "Add" (1) and is processed correctly but not accepted by the NAS (i.e., it does not pass the conditional access and admission control 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. When used in this situation, the Multicast Replication Control message:
  - MUST contain a Result set to 0x0;
  - MUST contain a Transaction ID with a unique value, as described in Section 3.6.1.6 of [RFC6320];
  - MUST contain the directive rejected by the NAS (i.e., Target TLV and Command TLV) but with a Command Code set to "Admission Control Reject" (4), "Conditional Access Reject" (5), or "Admission Control and Conditional Access Reject" (6) as applicable.
- 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 4.2 of [RFC6320]) 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:

- the Target TLV (Section 4.3 of [RFC6320] or an extension), specifying a single access line;
- 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 current 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 current 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 Result 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 Result Code field MUST be set to 0, 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.

- o 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 Result Code value set to 0x68 "Inconsistent views of delegated bandwidth amount" or 0x69 "Bandwidth request conflict" as applicable. (See below for more information).

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 Result 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 Result Code values applicable specifically to the contents of the Bandwidth-Request TLV. These Result 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.

0x67 Invalid preferred bandwidth amount.

Where detected: control application at the receiver of the Bandwidth Reallocation Request message.

Further description: the preferred and required amounts of bandwidth in the TLV do not have the numerical relationship described above.

Required additional information in the message: as described above.

Target: control application at the sender of the Bandwidth Reallocation Request message.

Action RECOMMENDED for the receiving ANCP agent: report the error to the control application with the returned value of the Bandwidth-Allocation TLV. See also Section 4.6.2.2.

0x68 Inconsistent views of delegated bandwidth amount.

Where detected: control application at the NAS.

Further description: the NAS has an outstanding Bandwidth Reallocation Request, so it is rejecting a similar request from the AN. In the AN request, the required amount was less than the NAS's view of the current amount of delegated bandwidth.

Required additional information in the message: as described above.

Target: control application at the AN.

Action RECOMMENDED for the receiving ANCP agent: report the error to the AN control application with the returned value of the Bandwidth-Allocation TLV. See also Section 4.6.2.2.

0x69 Bandwidth request conflict.

Where detected: control application at the NAS.

Further description: the NAS has an outstanding Bandwidth Reallocation Request, so it is rejecting a similar, valid request from the AN.

Required additional information in the message: as described above.

Target: control application at the AN.

Action RECOMMENDED for the receiving ANCP agent: report the error to the AN control application with the returned value of the Bandwidth-Allocation TLV. See also Section 4.6.2.2.

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 Result Code field equal to 0x68 "Inconsistent views of delegated bandwidth amount" or 0x69 "Bandwidth request conflict". If Result Code value 0x68 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 Result Code value 0x69 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 Result 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 Result Code value MUST be set to 0. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in Section 3.6.1.6 of [RFC6320].

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 Result Code field MUST be set to 0. 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 Result Code field MUST be set to the value that indicates the nature of the error (e.g., 0x500 "One or more of the specified ports do not exist"). 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 Result Code field MUST be set to 0x000. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in section 3.6.1.6 of [RFC6320].

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

- o the AN's view of which multicast flows are currently active on a specified set of access ports; or

- o the AN's view of the access ports on which a specified set of multicast flows are currently active; or

- o the AN's view of all the multicast flows currently active on each access 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 [RFC6320].

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 5.12.

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 Result Code field MUST be set to 0. 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:

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

- o 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 Result 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 5.14.

4.10.1. Sender Behaviour

The sender of a Committed Bandwidth Report message MUST set the Result field to Ignore (0x0). The Result Code field MUST be set to 0x000. The sender MUST populate the ANCP Transaction Identifier field with a unique value, as described in section 3.6.1.6 of [RFC6320].

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.

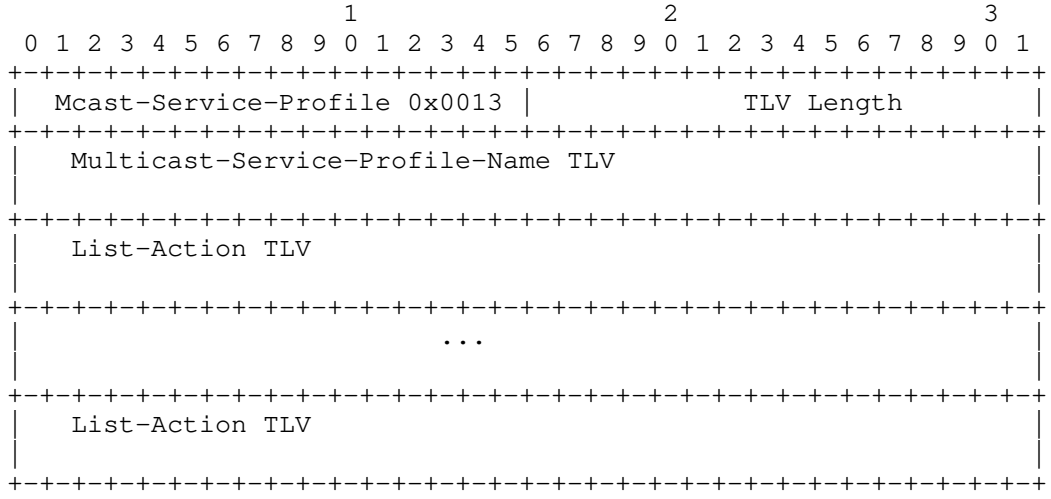

Figure 6: Multicast-Service-Profile TLV

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

- The Multicast-Service-Profile TLV Type is 0x0013.
- 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 which 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.

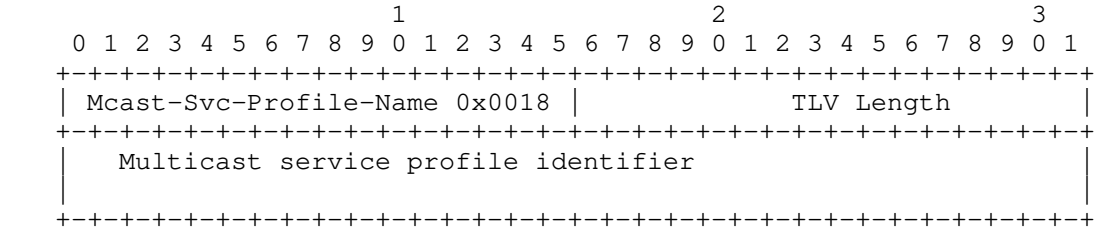

Figure 7: Multicast-Service-Profile-Name TLV

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

- The Multicast-Service-Profile-Name TLV Type is 0x0018.
- TLV Length: up to 255 octets.
- Multicast service profile identifier: an opaque sequence of octets 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.

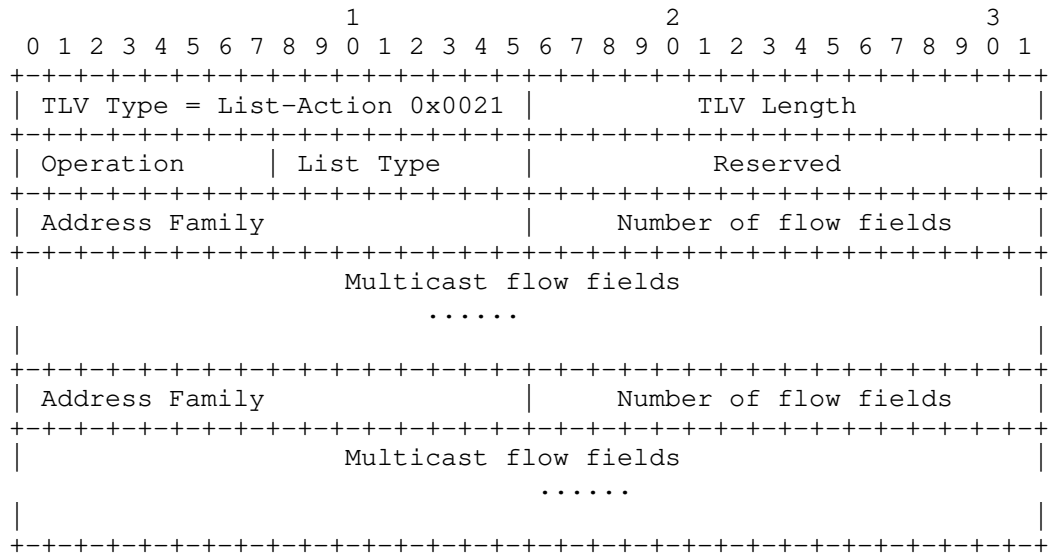

Figure 8: List-Action TLV

The List-Action TLV contains the following fields:

- The List-Action TLV Type is 0x0021.
- 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:
  - 1 "Add": the multicast flow fields are to be added to the list.
  - 2 "Delete": 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).
  - 3 "Replace": 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 1 "White", 2 "Black", or 3 "Grey".

- Reserved: a sender MUST set this field to zeroes. 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 1 "IPv4" or 2 "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:

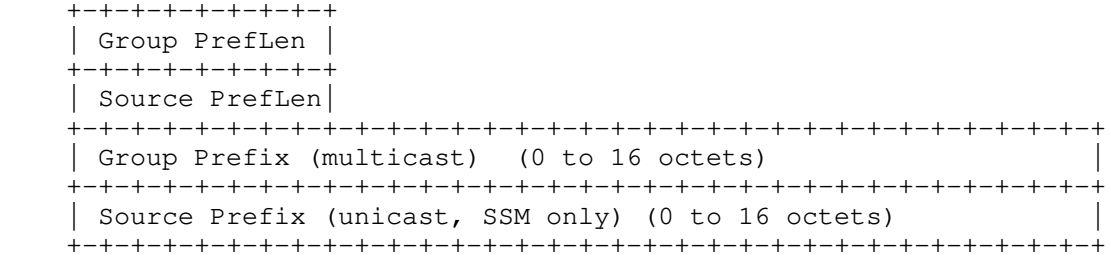

Figure 9: Organization of a Single Multicast Flow Field

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 0 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 0 is provided for a particular

prefix length, the corresponding prefix MUST be omitted from the field contents.

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 in a Status-Info TLV, in a Generic Response reporting a failed command in 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.

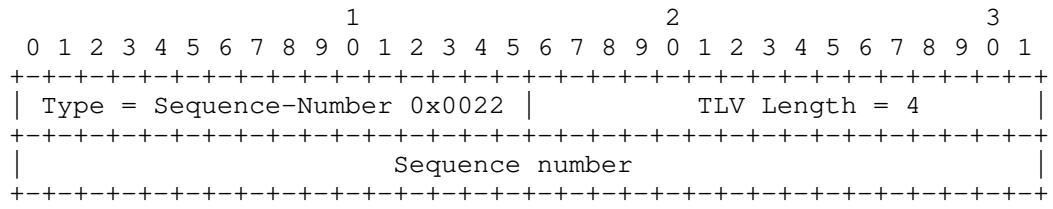

Figure 10: Sequence-Number TLV

The Sequence-Number TLV has the following fields:

- The Sequence-Number TLV Type is 0x0022.
- TLV length is 4.
- 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.

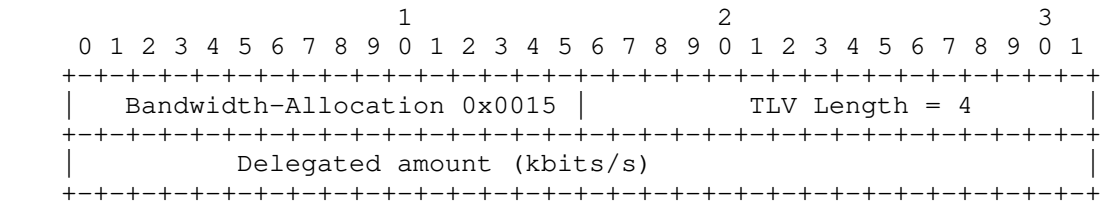

Figure 11: The Bandwidth-Allocation TLV

The Bandwidth-Allocation TLV has the following fields:

- The Bandwidth-Allocation TLV Type is 0x0015.
- 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.

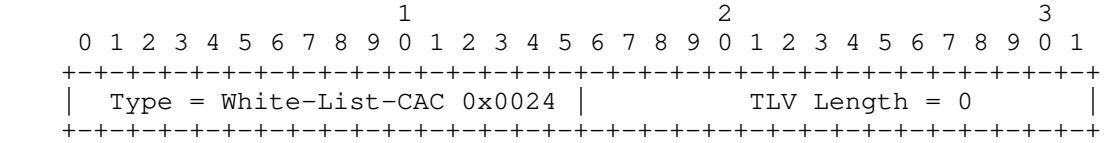

Figure 12: White-List-CAC TLV

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

- The White-List-CAC TLV Type is 0x0024.
- 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.

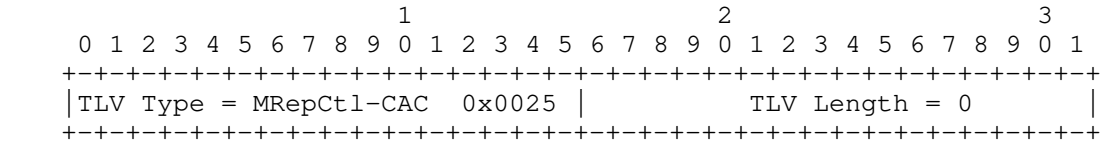

Figure 13: MRepCtl-CAC TLV

The MRepCtl-CAC TLV contains the following fields:

- The MRepCtl-CAC TLV Type is 0x0025.
- 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.

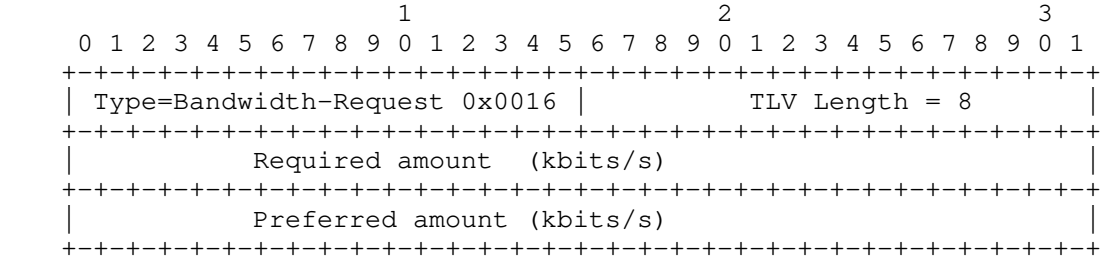

Figure 14: The Bandwidth-Request TLV

The Bandwidth-Request TLV has the following fields:

- The Bandwidth-Request TLV Type is 0x0016.
- 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.

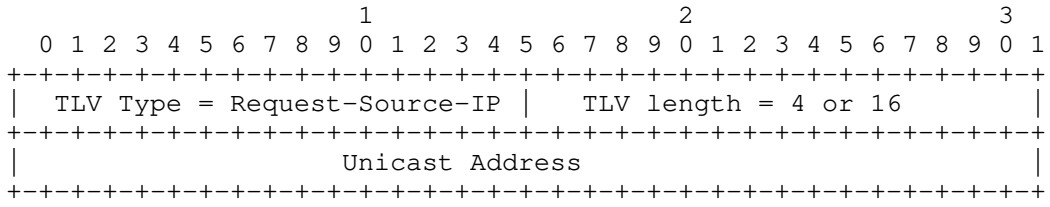

Figure 15: Request-Source-IP TLV

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

- The Request-Source-IP TLV Type is 0x0092.

- 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.

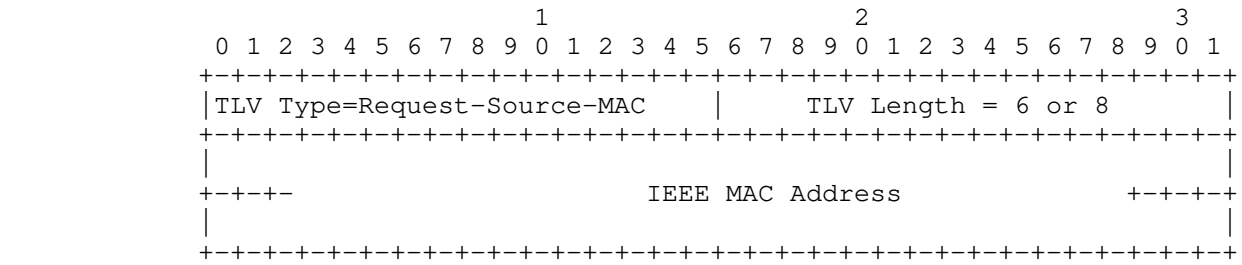

Figure 16: Request-Source-MAC TLV

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

- The Request-Source-MAC TLV Type is 0x0093.
- 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. Request-Source-Device-Id TLV

The Request-Source-Device-Id TLV provides a local identifier of the entity that originated a specific request to join or leave a multicast channel. The TLV is illustrated in Figure 17.

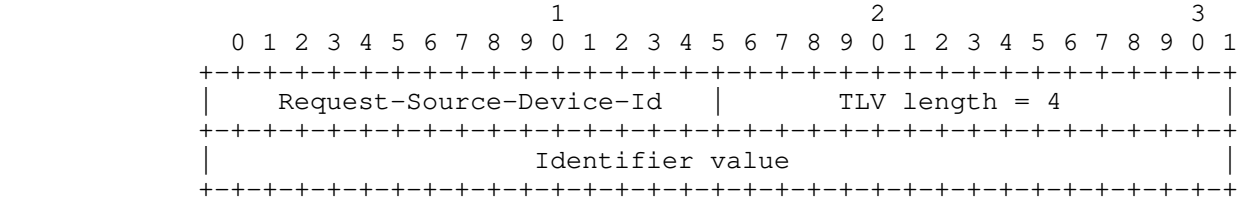

Figure 17: Request-Source-Device- Id TLV

The Request-Source-Device-Id TLV contains the following fields:

- The Request-Source-IP TLV Type is 0x0096.
- TLV length is 4.

- Local device identifier value, known to the AN and AAA. Given that the scope of the identifier is a single customer network, 32 bits is a more than sufficient numbering space.

5.12. Multicast-Flow TLV

IGMPv3 [RFC3376] and MLDv2 [RFC3810] allow multicast listeners to specify multiple source addresses for the same multicast group. Similarly the Multicast-Flow TLV specifies a multicast flow in terms of its multicast group address and, if applicable, one or more unicast source addresses. The Multicast-Flow TLV is illustrated in Figure 18.

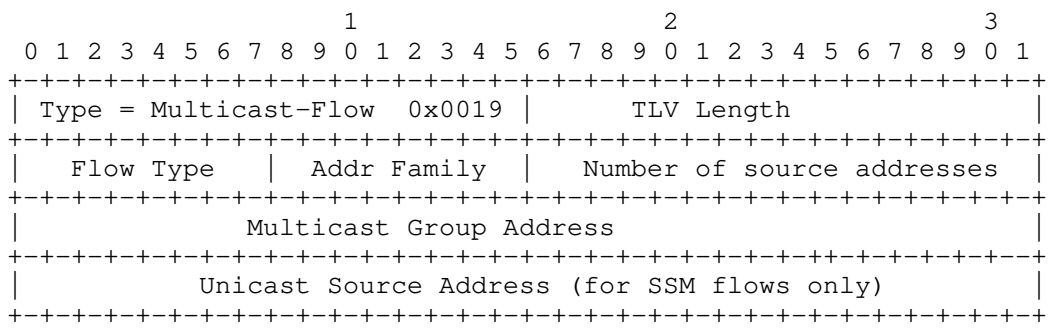

Figure 18: Multicast-Flow TLV

The Multicast-Flow TLV has the following fields:

- The Multicast-Flow TLV Type is 0x0019.
- TLV Length: ranges from a minimum of 8 (for an ASM IPv4 flow) upwards. Total length is 4 + 4*(Number of Source Addresses +1) for IPv4 or 4 + 16*(Number of Source Addresses + 1) for IPv6.
- Flow Type: 1 "Any Source Multicast (ASM)", 2 "Source-Specific Multicast (SSM)".
- Addr Family: address family of the multicast source and group addresses, encoded in accordance with the IANA PIM Address Family registry ([PIMreg]). 1 indicates IPv4, 2 indicates IPv6.
- Number of Source Addresses: 0 for ASM, 1 or more for SSM.
- Multicast Group Address: a multicast group address within the given address family. The group address MUST always be present.

- o Unicast Source Address: unicast address within the given address family. If the Flow Type is "ASM" (1), a source address MUST NOT be present. If the Flow Type is "SSM" (2), the number of source addresses given by the Number of Source Addresses field MUST be present.

The full versions of IGMPv3 and MLDv2 support both INCLUDE and EXCLUDE modes for specifying the desired sources for SSM flows. The Multicast-Flow TLV supports INCLUDE mode only. [RFC5790] (Lightweight IGMPv3 and MLDv2) provides guidance on converting EXCLUDE mode IGMP/MLD records to INCLUDE mode for the Multicast-Flow TLV.

5.13. 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 19.

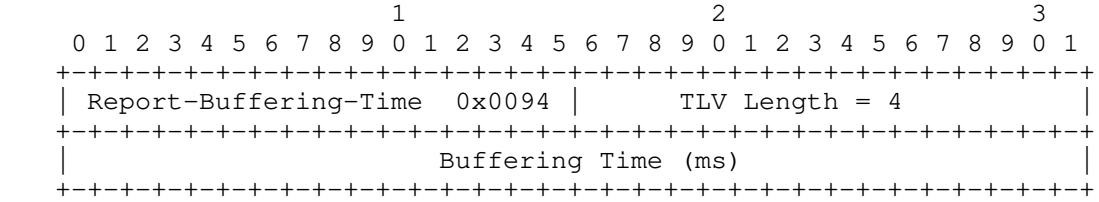

Figure 19: Report-Buffering-Time TLV

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

- o The Report-Buffering-Time TLV Type is 0x0094.
- o TLV length is 4 octets.
- o Buffering Time is a 32-bit unsigned integer containing a time value in ms.

5.14. 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 20.

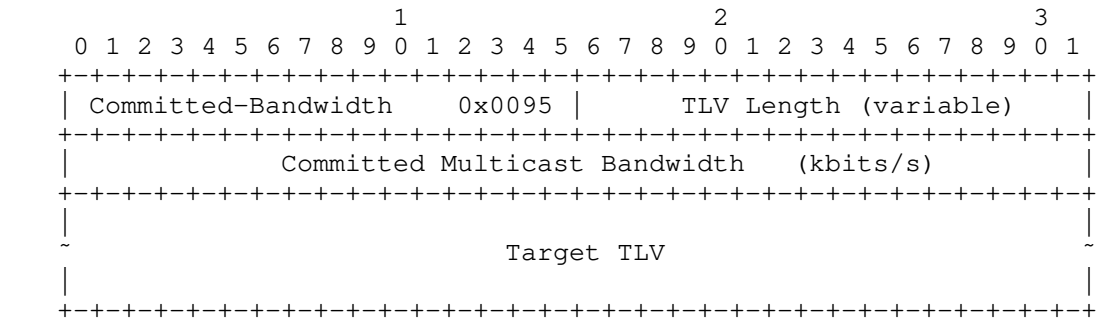

Figure 20: Committed-Bandwidth TLV

The Committed-Bandwidth TLV contains the following fields:

- The Committed-Bandwidth TLV Type is 0x0095.
- 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 3.5 of [RFC6320] 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 3);
- committed multicast bandwidth reporting (capability type 5);
- conditional access and admission control with white and black lists (capability type 6);
- conditional access and admission control with grey lists (capability type 7);
- bandwidth delegation (capability type 8).

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.

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.

In addition to the elements listed below, implementation of the Target TLV (Section 4.3 of [RFC6320]) is REQUIRED for all of the capabilities specified in this document.

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. Additionally, implementation of the Multicast Replication Control message requires implementation of the Command TLV (Section 4.4 of [RFC6320] with additional details in Section 4.3 of this document).

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	Sequence Number TLV;
Section 5.5	Bandwidth-Allocation TLV;
Section 5.7	MRepCtl-CAC TLV;
Section 5.12	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 committed multicast bandwidth reporting 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 5.13	Report-Buffering-Timer TLV;
Section 5.14	Committed-Bandwidth TLV;
Section 5.12	Multicast-Flow TLV.

Table 2: Protocol Requirements For Committed Multicast Bandwidth Reporting

6.1.3. Protocol Requirements For Conditional Access and Admission Control 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 and admission control 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 5.12	Multicast-Flow TLV.

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

6.1.4. Protocol Requirements For Conditional Access and Admission Control With Grey Lists

Table 4 specifies the protocol elements within Section 4 and Section 5 that MUST be implemented to support the conditional access and admission control with grey lists multicast capability. Additionally, implementation of the Multicast Replication Control message requires implementation of the Command TLV (Section 4.4 of [RFC6320] with additional details in Section 4.3 of this document).

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	Sequence 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	Request-Source-Device-Id TLV;
Section 5.12	Multicast-Flow TLV.

Table 4: Protocol Requirements For Conditional Access and Admission Control 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 5.12	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 corresponding 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 and Admission Control 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 group 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 and Admission Control 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 and admission control 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 requesting device.

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 that 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 the AN receives 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 and Admission Control With White and Black Lists and Conditional Access and Admission Control 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.

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.

> The procedures in the previous paragraphs imply that the AN has to retain a memory of whether an admitted flow was white-listed or grey-listed at the time of its admission/readmission.

6.3.2. Combination of Conditional Access and Admission Control 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. Conditional access follows the rules given in those sections in terms of matching flows against white and black and/or grey lists. When admission control is enabled at the AN, the amount of bandwidth used by the AN 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, if the AN is to do admission control for flows added by Multicast Replication Control messages, regardless of whether they are part of NAS-initiated replication or grey list multicast service processing, the NAS includes the MRepCtl-CAC TLV in a Provisioning message and the Bandwidth-Allocation TLV in a Port Management message. If instead the NAS will do admission control for flows added by Multicast Replication Control messages, regardless of whether they are part of NAS-initiated replication or grey list multicast service processing, it does not send the MRepCtl-CAC TLV in a Provisioning messages to the AN. The NAS can independently enable admission control for white flows on the AN by including 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 place 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 and admission control with grey lists.

7. Miscellaneous Considerations

This section deals with two sets of considerations. "Report Buffering Considerations" considers requirements for configuration in support of some of the Committed Bandwidth Reporting capability. "Congestion Considerations" is a warning to implementors about the possibility of control plane congestion, with suggestions for mitigation.

7.1. Report Buffering Considerations

The Committed Bandwidth Reporting capability allows the provisioning of a report buffering period to reduce the number of messages the AN passes to the NAS. An appropriate value for this period, if buffering is allowed at all, depends first on the effect of delay in reporting bandwidth changes, and secondly on the rate at which bandwidth changes are expected to occur.

Let us assume in the first instance that a delay in adjusting hierarchical scheduling at the NAS causes additional bandwidth demand to be served momentarily on a best-effort basis, introducing the possibility of jitter and, more crucially, packet loss. ITU-T Recommendation G.1080 [ITU-T_G.1080] Appendix IV indicates that the maximum tolerable duration of a loss episode is less than 16 ms. This would more likely apply in the middle of a programme rather than when it was starting up, but at least gives an (extremely conservative) order of magnitude for setting the buffering period.

The next question is whether enough messaging is likely to be generated that multiple bandwidth changes would be observed within such an interval. Let us consider a reasonable example in a DSL environment, where during the busiest hour of the day subscribers start watching at the rate of one programme per subscriber per hour. Typically, because of programme scheduling, the new channel requests might be concentrated within a three-minute period, giving an effective request rate of 1/(3 minutes * 60 seconds * 1000 ms/second) * 16 ms = 0.00009 requests per buffering interval of 16 ms. With these figures, an AN serving 10,000 subscribers will report an average of 0.9 bandwidth changes per 16 ms buffering interval. It appears that buffering is worthwhile only for larger-scale deployments.

Note that simple replacement of one channel with another -- channel surfing -- does not require reporting or adjustment at the NAS end.

7.2. Congestion Considerations

Implementors must beware of the possibility that a single channel-surfing subscriber could generate enough control messaging to overload the AN or the messaging channel between the AN and the NAS. The implementation problem is to strike the right balance between minimizing the processing of requests that have been overtaken by subsequent events and meeting requirements for what is termed "channel zapping delay". Nominally such a requirement is to be found in [ITU-T_G.1080] Section 8.1, but unfortunately no quantitative value was available at the time of publication of this document. Implementors will therefore have to base their work on discussions with customers until standardized requirements do become available. (It is possible that regional bodies or more specialized bodies have overtaken the ITU-T in this regard.)

A typical strategy for minimizing the work associated with request processing includes deliberate buffering of Join requests for a short period in case matching Release requests are detected, followed by discard of both requests. More generally, processing of requests from individual subscribers may be rate limited, and the global rate of messaging to the NAS can also be limited. If the AN gets overloaded, deliberate dropping of stale requests can be implemented, for some definition of "stale".

8. Security Considerations

The security considerations of ANCP are discussed in [RFC6320] and in [RFC5713]. Multicast does not in principle introduce any new security considerations, although it does increase the attractiveness of the ANCP protocol as a means of denial of service (e.g., through direction of multicast streams onto the target) or theft of service.

As mentioned in Section 4.4, the inclusion of the Request-Source-MAC or Request-Source-IP TLV in the Multicast Admission Control message presents privacy issues. An attacker able to get access to the contents of this message would, like the content provider, be able to track consumption of multicast content to the individual device and potentially to individual persons if they are associated with particular devices. To make the connection between devices and individuals, the attacker needs to get information from sources other than ANCP, of course, but let us assume that this has happened.

The protection specified for ANCP in [RFC6320] will apply to the transmission of the Multicast Admission Control message across the access network to the NAS. Hence the attacker’s potential points of access are between the subscriber and the AN, at the AN and at the NAS. Moreover, if the MAC or IP address are transmitted onwards from

the NAS to AAA in a request for policy, that whole onward path has to be examined for vulnerability.

The question is how many of these potential points of attack can be eliminated through operational practice. The segment from the subscriber through the AN itself seems out of scope of this discussion -- protection of this segment is basic to subscriber privacy in any event, and likely a business requirement. The segment from the AN to the NAS is covered by the basic ANCP protection specified in RFC 6320. This leaves the NAS and the path between the NAS and AAA for consideration.

The operator can eliminate the path between the NAS and AAA as a point where the attacker can access per-device information by downloading per-device policy to the NAS for all identified user devices for the particular subscriber. The NAS then selects the applicable policy based on the particular device identifier it has received. This is as opposed to the NAS sending the identifier of the device in question to AAA and getting policy just for that device.

The alternative is to protect the path between the NAS and AAA. If Diameter is used as the AAA protocol, Section 2.2 of [RFC6733] mandates use of IPsec, TLS/TCP, or DTLS/SCTP for that purpose. If RADIUS is used, the operator should deploy TLS transport as specified in [RFC6614].

This leaves the NAS itself as a point of attack. In theory the NAS could be eliminated if the AN remapped the requesting MAC or IP address to an identifier known to itself and AAA, but not the NAS. This would require local configuration on the AN, which may be possible under some circumstances. The Request-Source-Device-Id TLV specified in Section 5.11 is available to transmit such an identifier in place of the Request-Source-MAC or Request-Source-IP.

9. IANA Considerations

IANA NOTE: Please replace XXXX with the RFC number of this document.

This document defines the following additional values within the ANCP Message Type Name Space registry:

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

This document defines the following additional values for the ANCP Result Code registry. In support of these assignments, IANA is requested to change the lower limit of 0x100 specified by [RFC6320] for assignments by IETF Consensus to 0x64.

Result Code	One-Line Description	Reference
0x64	Command error.	RFC XXXX
0x65	Invalid flow address.	RFC XXXX
0x66	Multicast flow does not exist.	RFC XXXX
0x67	Invalid preferred bandwidth amount.	RFC XXXX
0x68	Inconsistent views of delegated bandwidth amount.	RFC XXXX
0x69	Bandwidth request conflict.	RFC XXXX

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

Command Code Value	Command Code Directive Name	Reference
1	Add	RFC XXXX
2	Delete	RFC XXXX
3	Delete All	RFC XXXX
4	Admission Control Reject	RFC XXXX
5	Conditional Access Reject	RFC XXXX
6	Admission Control and Conditional Access Reject	RFC XXXX

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

Type Code	TLV Name	Reference
0x0013	Multicast-Service-Profile	RFC XXXX
0x0015	Bandwidth-Allocation	RFC XXXX
0x0016	Bandwidth-Request	RFC XXXX
0x0018	Multicast-Service-Profile-Name	RFC XXXX
0x0019	Multicast-Flow	RFC XXXX
0x0021	List-Action	RFC XXXX
0x0022	Sequence-Number	RFC XXXX
0x0024	White-List-CAC	RFC XXXX
0x0025	MRepCtl-CAC	RFC XXXX
0x0092	Request-Source-IP	RFC XXXX
0x0093	Request-Source-MAC	RFC XXXX
0x0094	Report-Buffering-Time	RFC XXXX
0x0095	Committed-Bandwidth	RFC XXXX
0x0096	Request-Source-Device-Id	RFC XXXX

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

Value	Capability Type Name	Tech Type	Capability Data?	Reference
3	NAS-Initiated Replication	0	No	RFC XXXX
5	Committed Bandwidth Reporting	0	No	RFC XXXX
6	Conditional Access With White and Black Lists	0	No	RFC XXXX
7	Conditional Access With Grey Lists	0	No	RFC XXXX
8	Bandwidth Delegation	0	No	RFC XXXX

10. 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.

Thanks to the Document Shepherd, Matthew Bocci, and Area Director, Ted Lemon, for points raised by their reviews following Working Group Last Call.

Further thanks to Dacheng Zhang, Mehmet Ersue, and Christer Holmberg for their reviews on behalf of the Security, Operations, and Gen-Art directorates. Dacheng's comments led to changes at several points in the draft, while Mehmet's led to creation of the Miscellaneous Considerations section. Finally, thanks to Brian Haberman for stimulating a review of the architectural assumptions and their relationship to the ability of user devices to obtain access to non-IPTV multicast services. This also led to changes in the draft.

11. References

11.1. Normative References

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

[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.

[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.

[RFC5790] Liu, H., Cao, W., and H. Asaeda, "Lightweight Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Version 2 (MLDv2) Protocols", RFC 5790, February 2010.

[RFC6320] Wadhwa, S., Moisand, J., Haag, T., Voigt, N., and T. Taylor, "Protocol for Access Node Control Mechanism in Broadband Networks", RFC 6320, October 2011.

11.2. Informative References

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

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

[ITU-T_G.1080] ITU-T, "ITU-T Recommendation G.1080: Quality of experience requirements for IPTV services", December 2008.

[RFC2236] Fenner, W., "Internet Group Management Protocol, Version 2", RFC 2236, November 1997.

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

[RFC5713] Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", RFC 5713, January 2010.

[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.

[RFC6614] Winter, S., McCauley, M., Venaas, S., and K. Wierenga, "Transport Layer Security (TLS) Encryption for RADIUS", RFC 6614, May 2012.

[RFC6733] Fajardo, V., Arkko, J., Loughney, J., and G. Zorn, "Diameter Base Protocol", RFC 6733, October 2012.

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 21. The message has the following contents:

- Message type is 93.
- The Result and Result Code fields in the header are set to zeroes, as specified [RFC6320].
- 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.

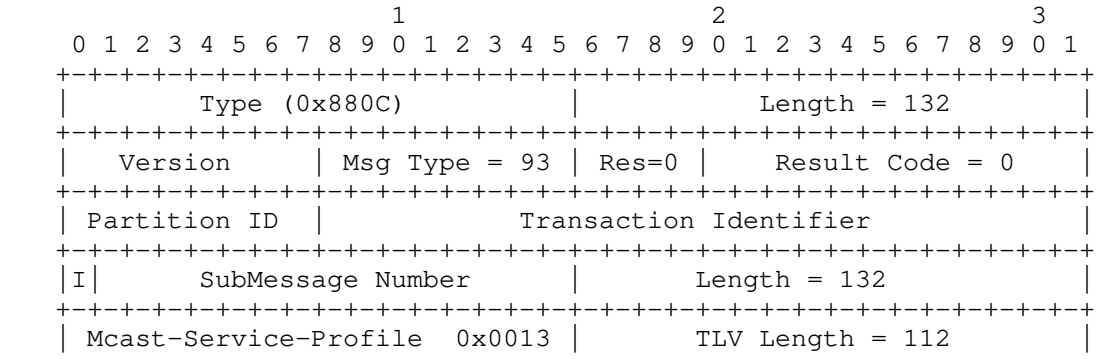

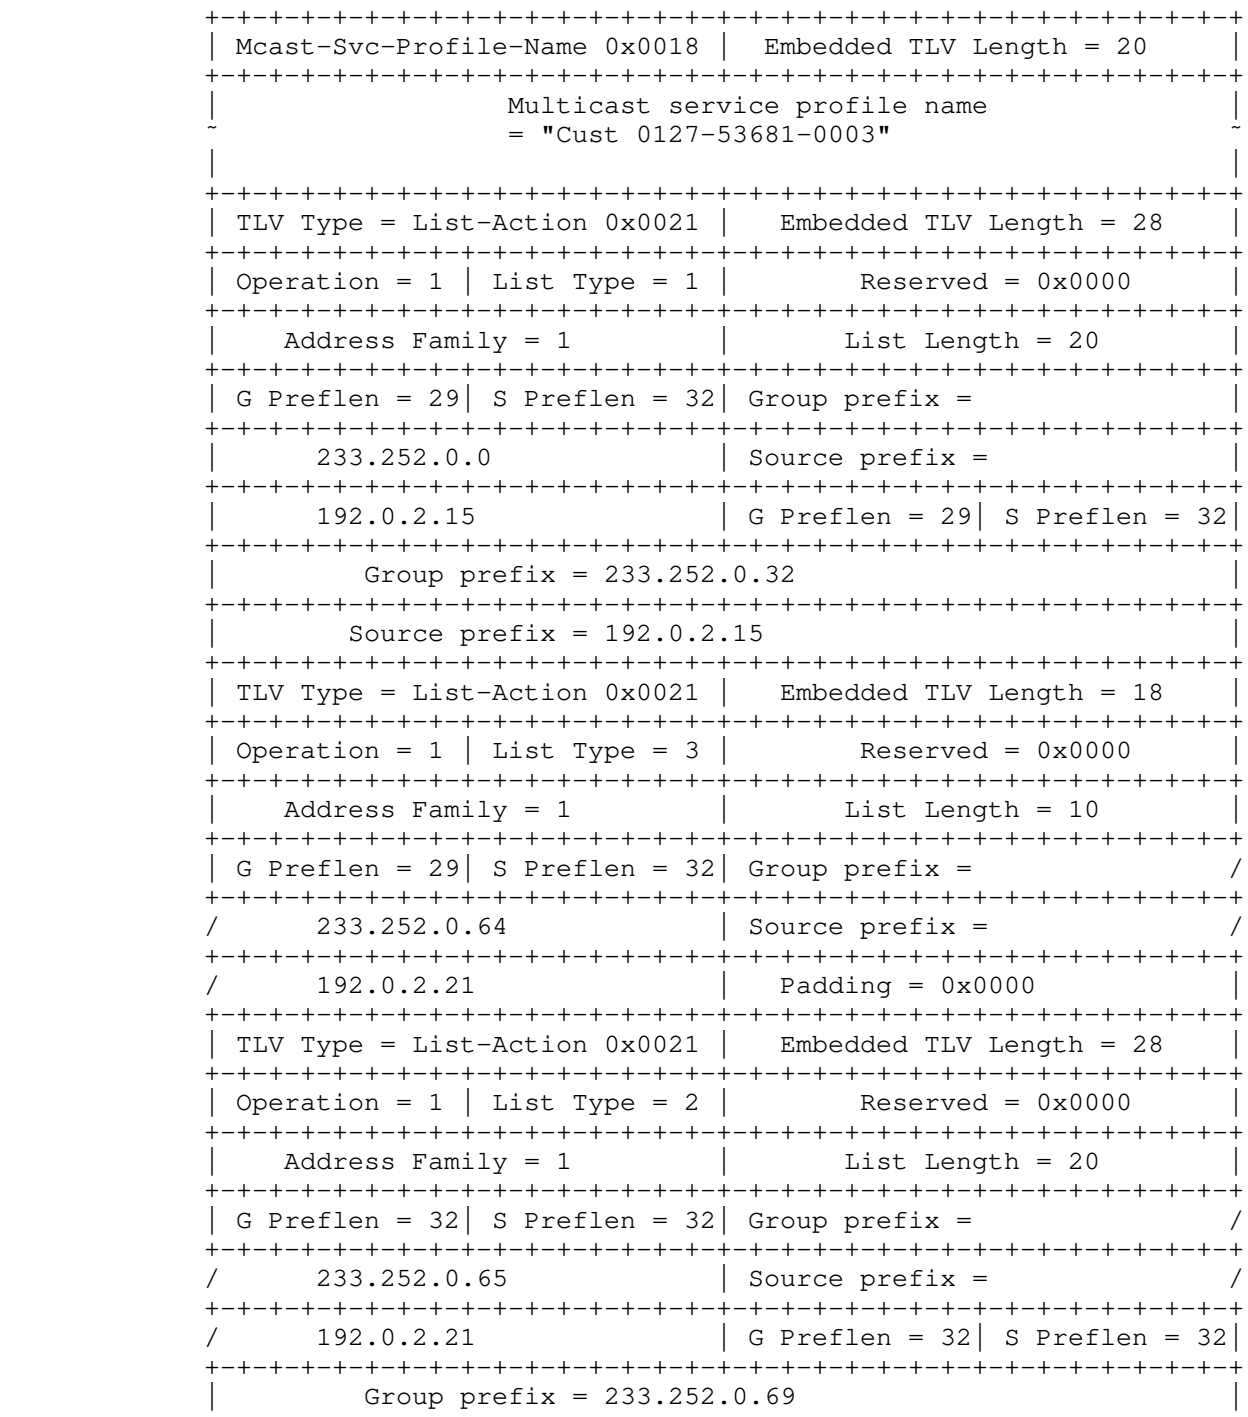
Mcast-Svc-Profile-Name 0x0018
Embedded TLV Length = 20
Multicast service profile name
= "Cust 0127-53681-0003"
TLV Type = List-Action 0x0021
Embedded TLV Length = 28
Operation = 1
List Type = 1
Reserved = 0x0000
Address Family = 1
List Length = 20
G Preflen = 29
S Preflen = 32
Group prefix =
233.252.0.0
Source prefix =
192.0.2.15
G Preflen = 29
S Preflen = 32
Group prefix = 233.252.0.32
Source prefix = 192.0.2.15
TLV Type = List-Action 0x0021
Embedded TLV Length = 18
Operation = 1
List Type = 3
Reserved = 0x0000
Address Family = 1
List Length = 10
G Preflen = 29
S Preflen = 32
Group prefix =
233.252.0.64
Source prefix =
192.0.2.21
Padding = 0x0000
TLV Type = List-Action 0x0021
Embedded TLV Length = 28
Operation = 1
List Type = 2
Reserved = 0x0000
Address Family = 1
List Length = 20
G Preflen = 32
S Preflen = 32
Group prefix =
233.252.0.65
Source prefix =
192.0.2.21
G Preflen = 32
S Preflen = 32
Group prefix = 233.252.0.69

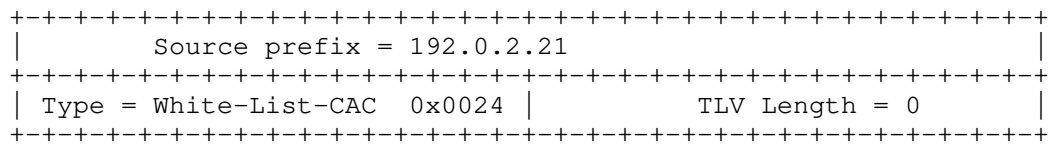

Figure 21: Example Provisioning Message

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 [RFC6320]. 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.

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 22.

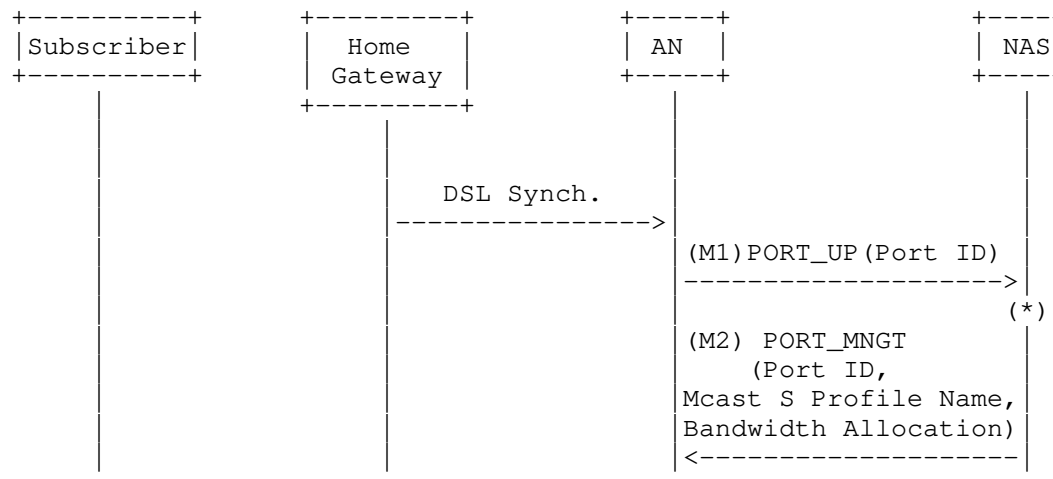

(*) The NAS may optionally seek direction from an external Autorization/Policy Server

Figure 22: 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 23) 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 (0x1) for this example.
- Result Code is 0.
- A transaction identifier is assigned by the NAS.
- 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 "Configure Connection Service Data" (8).

- X-Function is set to 0.

- Tech Type is "DSL" (5).

- 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.

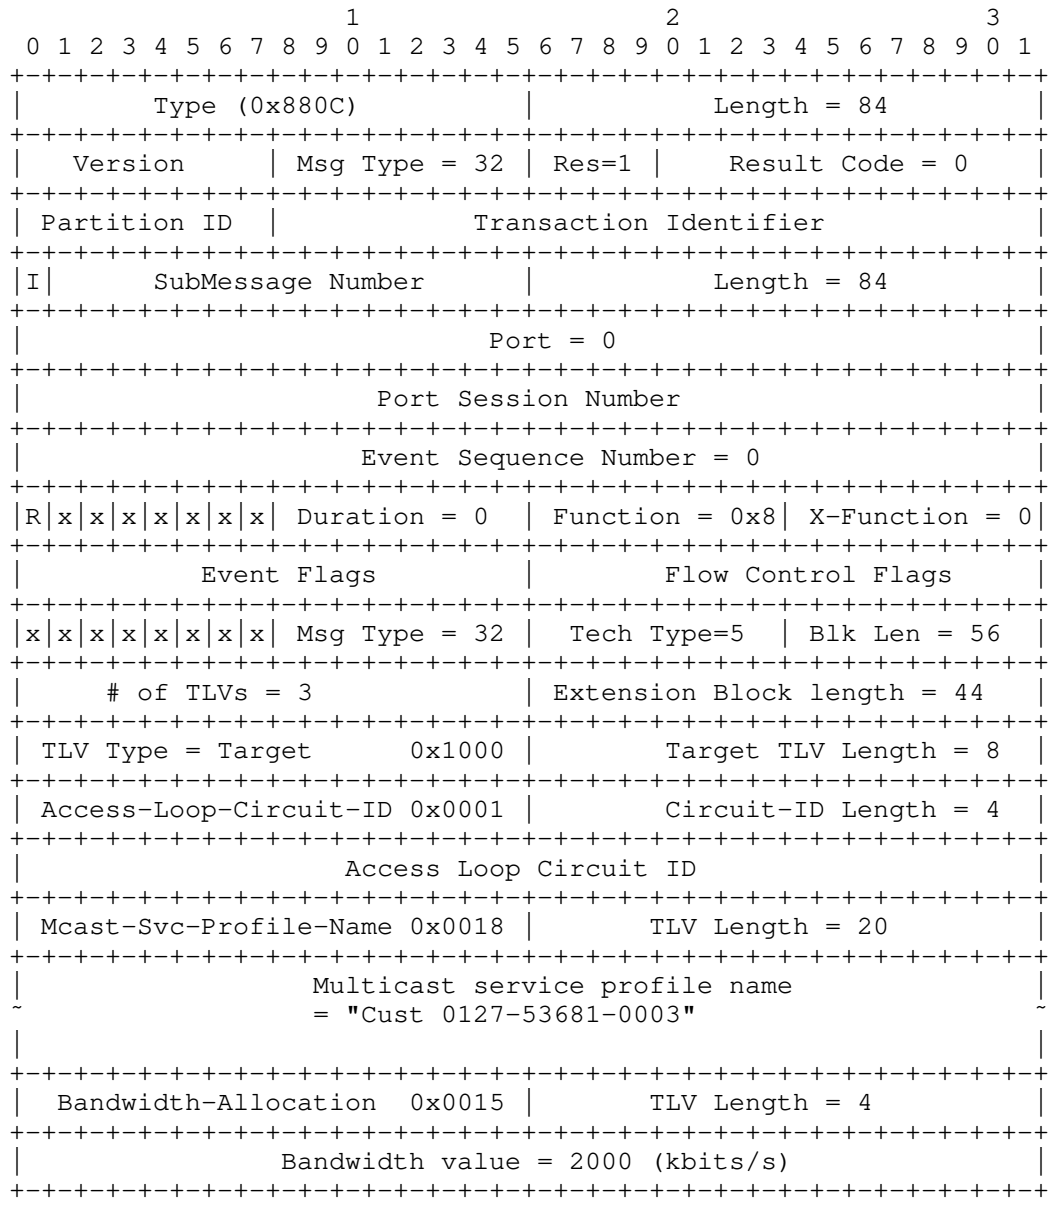

Figure 23: Example Port Management Message

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 24 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.

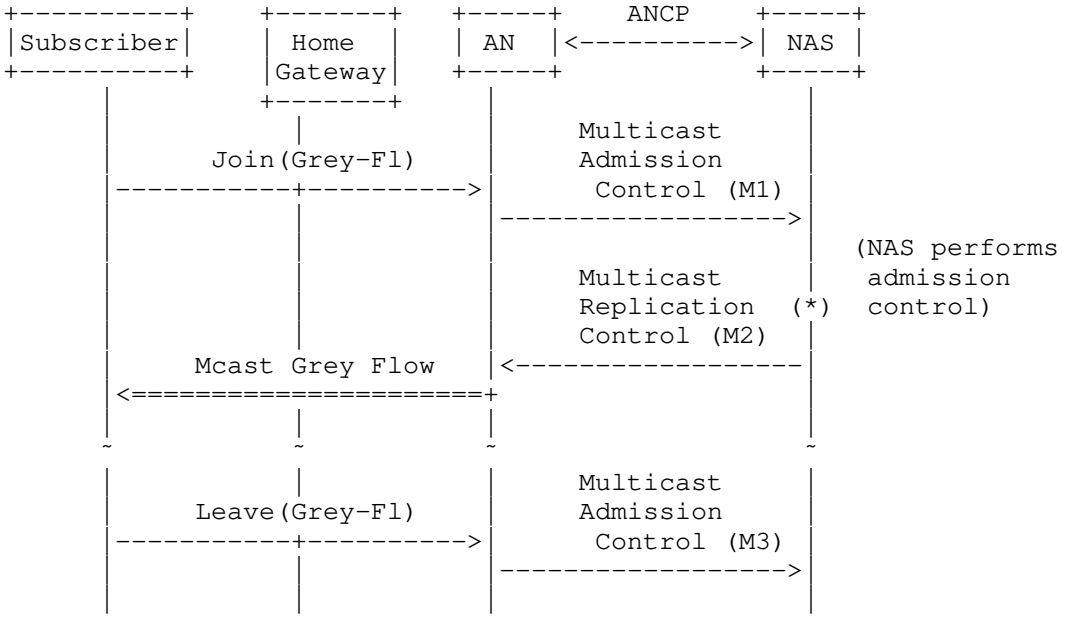

Grey-Fl : Multicast Flow matching an entry in grey List

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

Figure 24: Successful Join/Leave Operations, Grey-Listed Flow

The Multicast Admission Control message M1 contains:

- an ANCP Header with:
  - Message Type is 145;
  - Result = Ignore (0x0);
  - a transaction identifier assigned by the AN.
- a Target TLV identifying the AN Port
- a Command TLV containing:
  - Command Code = "Add" (1);
  - Accounting = "No" (0);
  - a Multicast-Flow embedded TLV indicating the multicast flow for which the AN received the IGMP Join: flow type "SSM" (2), address family "IPv4" (1), Group address = 233.252.0.67, Source Address = 192.0.2.21;
  - a Request-Source-Device-Id embedded TLV containing the IGMP join source local device identifier value 5.

The Multicast Admission Control message M1 is illustrated in Figure 25:

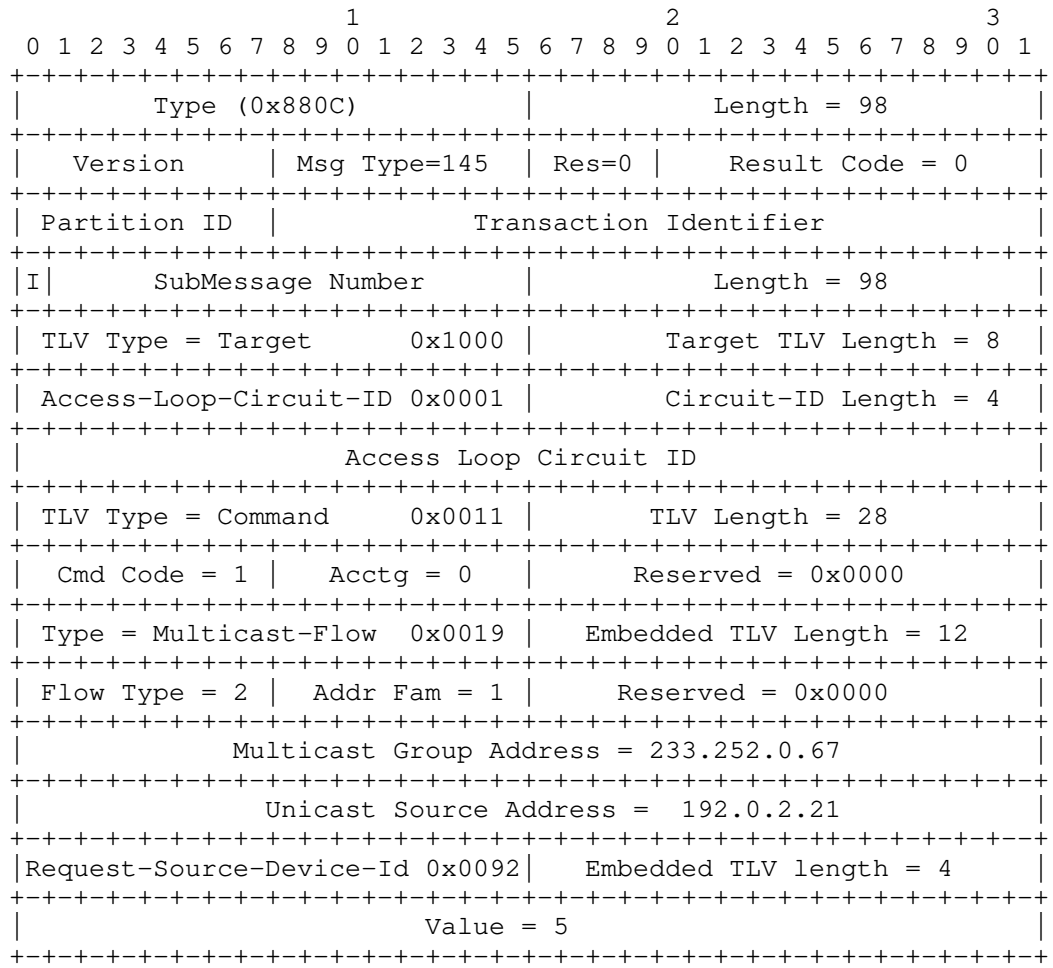

Figure 25: Multicast Admission Control Message Seeking To Add A Flow

The Multicast Replication Control message M2 contains:

- an ANCP Header with:
  - Message Type = "Multicast Replication Control" (144);
  - Result= 0x1 (NAck);
  - a transaction identifier assigned by the NAS;
- a Target TLV identifying the AN Port;

- a Command TLV containing:
  - Command Code = "Add" (1);
  - Accounting = "Yes" (1), since in our example the operator wants accounting on this flow.
  - a Multicast-Flow embedded TLV indicating the multicast flow that the NAS is admitting for this access line: flow type "SSM" (2), address family "IPv4" (1), Group address = 233.252.0.67, Source Address = 192.0.2.21.

The Multicast Admission Control message M2 is illustrated in Figure 26.

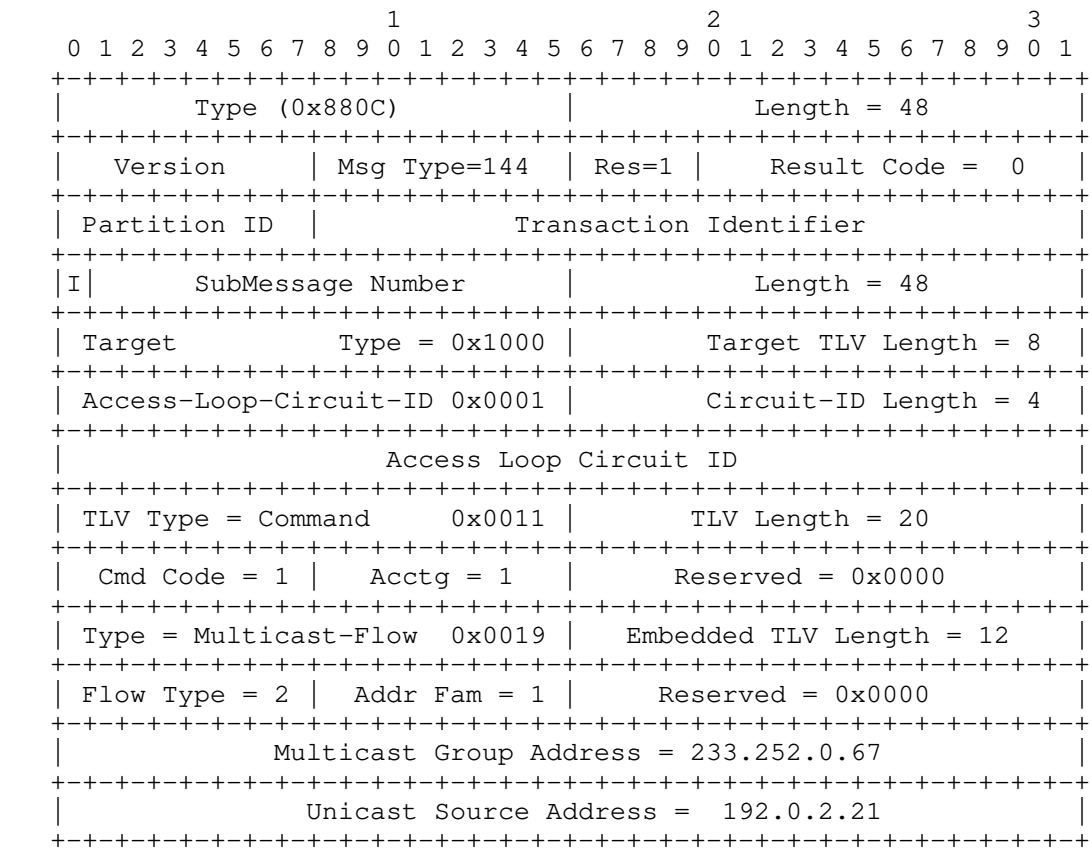

Figure 26: Multicast Replication Control Message Admitting A Flow

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

- an ANCP Header with:
  - Message Type is 145;
  - Result = Ignore (0x0)
  - a transaction identifier assigned by the AN;
- a Target TLV identifying the access line;
- a Command TLV containing:
  - a Command Code = "Delete" (2);
  - Accounting = "No" (0);
  - a Multicast-Flow embedded TLV indicating the multicast flow for which the AN received the IGMP leave: flow type "SSM" (2), address family "IPv4" (1), Group address = 233.252.0.67, Source Address = 192.0.2.21.
  - a Request-Source-Device-Id embedded TLV containing the IGMP leave request source, the device identified by the local value 5.

The Multicast Admission Control message M3 is illustrated in Figure 27.

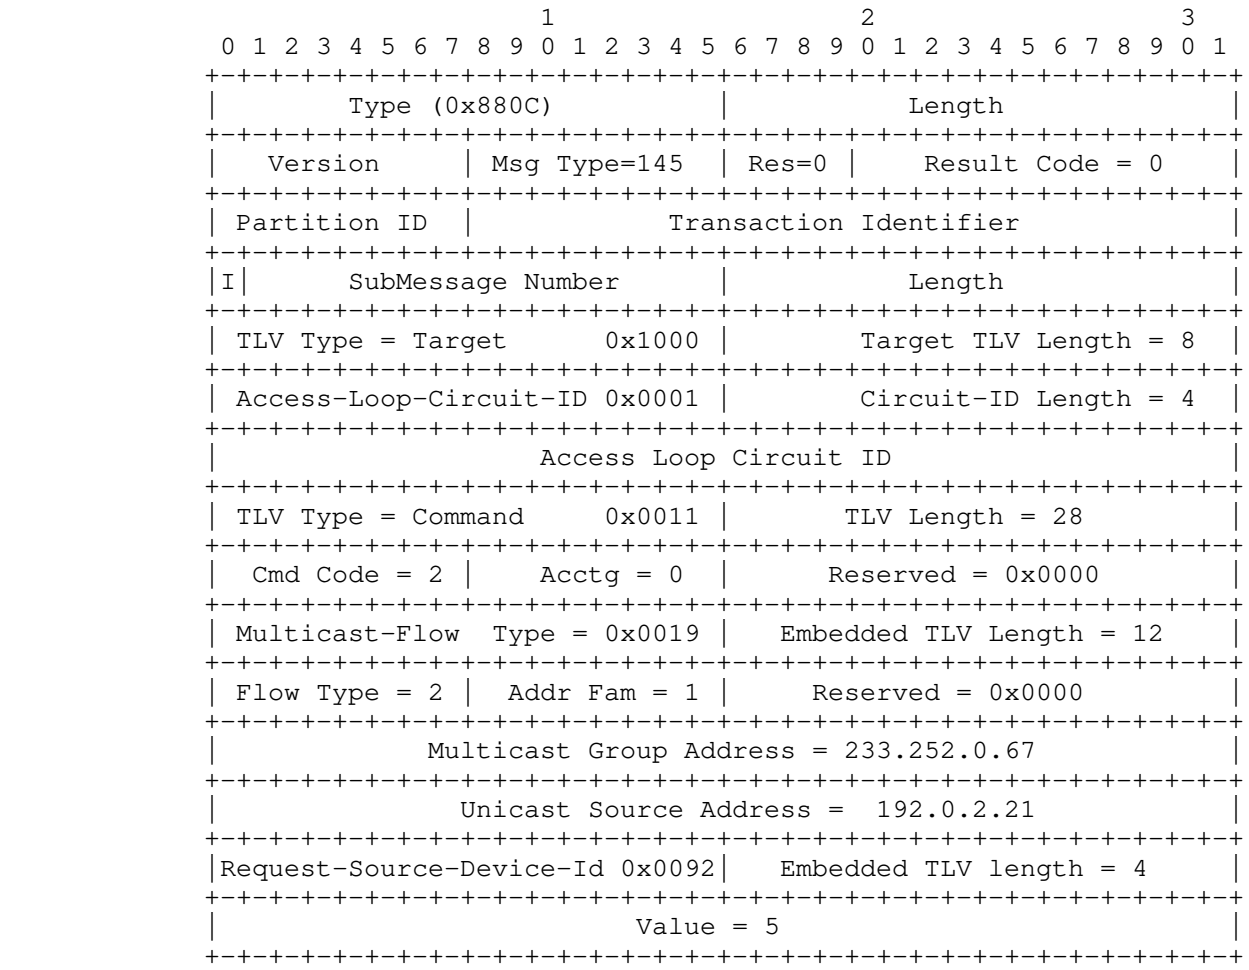

Figure 27: Multicast Admission Control Message Signalling Flow Termination

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 28.

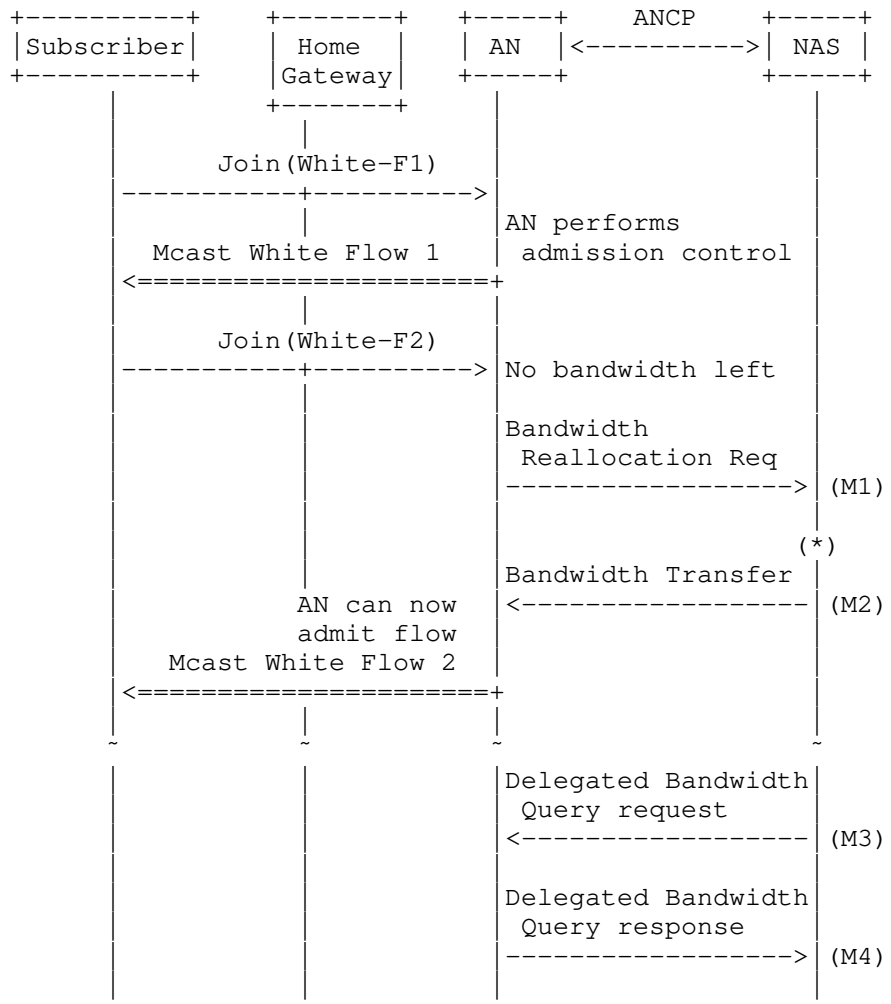

(*) The NAS may optionally seek direction from an external Authorization/Policy Server

Figure 28: Successful Join/Leave Operations, White-Listed Flow

The Bandwidth Reallocation Request message (M1) is shown in Figure 29. The contents require little explanation. The Message Type for the Bandwidth Reallocation Request is 146. The Result field is set to Ignore (0x0). Besides the Target, the message has one other TLV, the Bandwidth-Request, with a TLV Type of 0x0016. The TLV contains Required Amount and Preferred Amount fields, set to 4000 and 6000 kbits/s respectively.

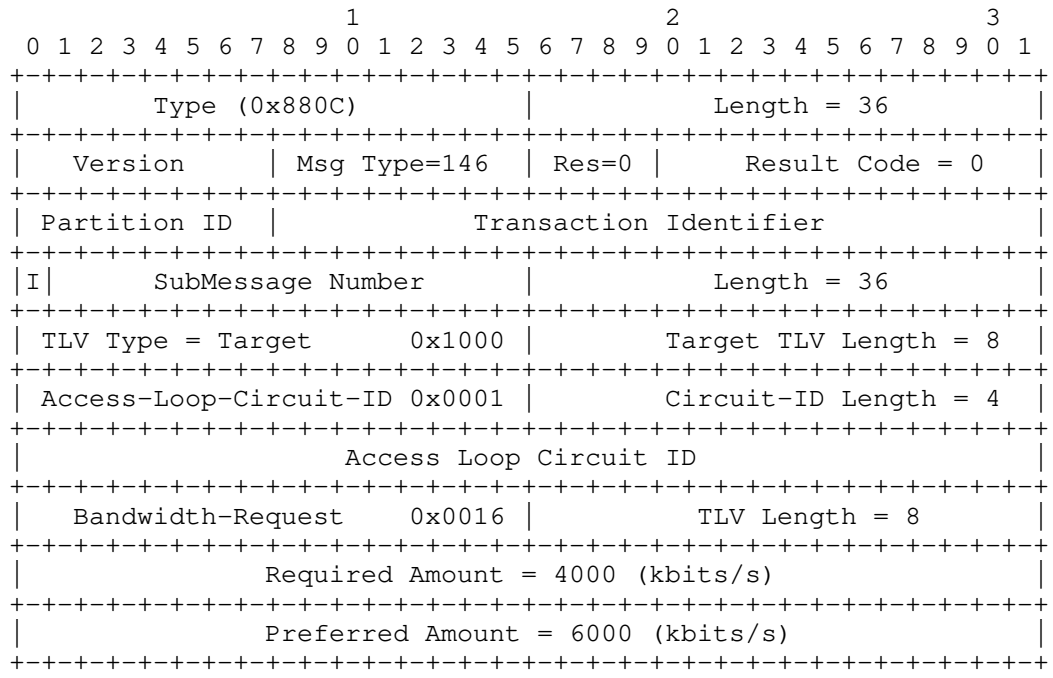

Figure 29: Bandwidth Reallocation Request Message

The Bandwidth Transfer message (M2) is shown in Figure 30. 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 0x0015 and contains a Delegated Amount field, set to 4000 kbits/s.

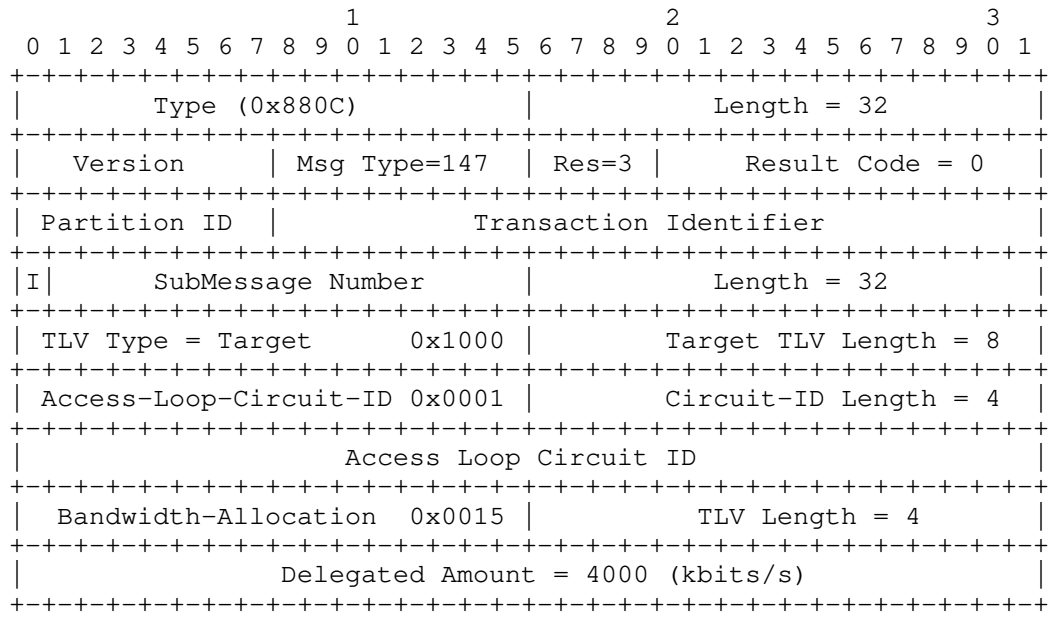

Figure 30: NAS Response, Bandwidth Transfer Message

The Delegated Bandwidth Query request message (M3) is shown in Figure 31. 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.

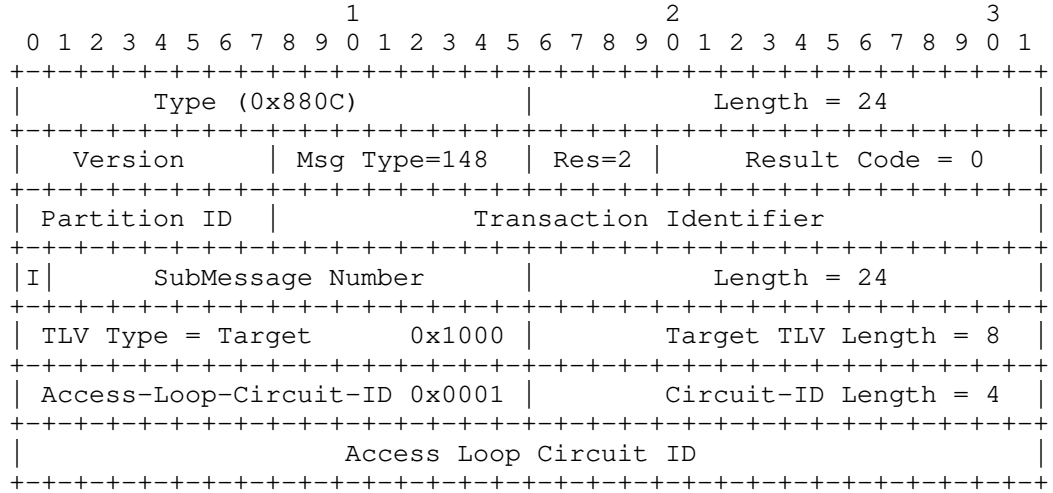

Figure 31: Delegated Bandwidth Query Request Message

Finally, the Delegated Bandwidth Query response message (M4) is shown in Figure 32. 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.

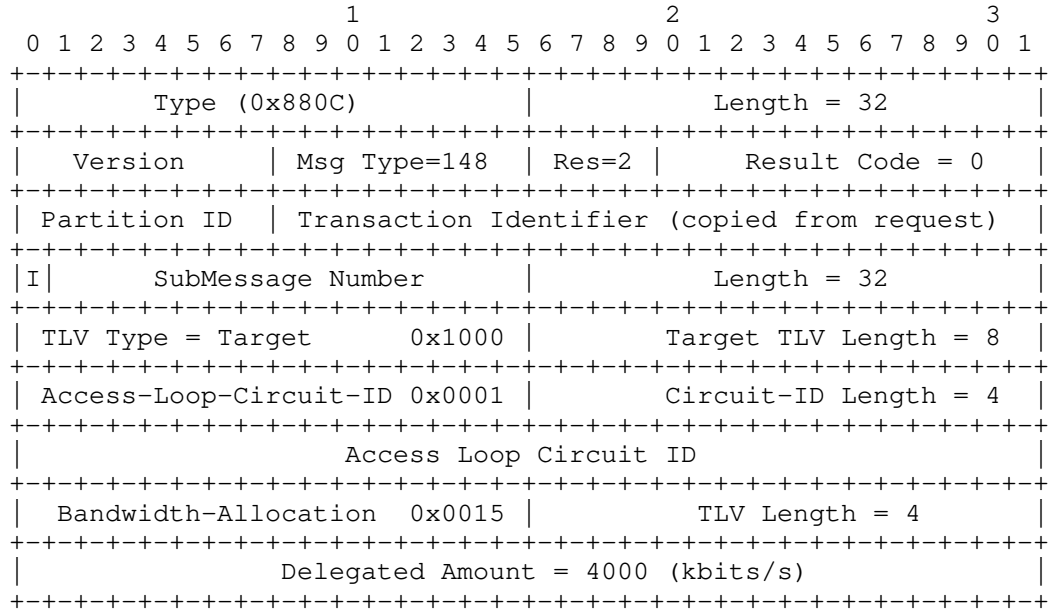

Figure 32: Delegated Bandwidth Query Response Message

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 33.

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.

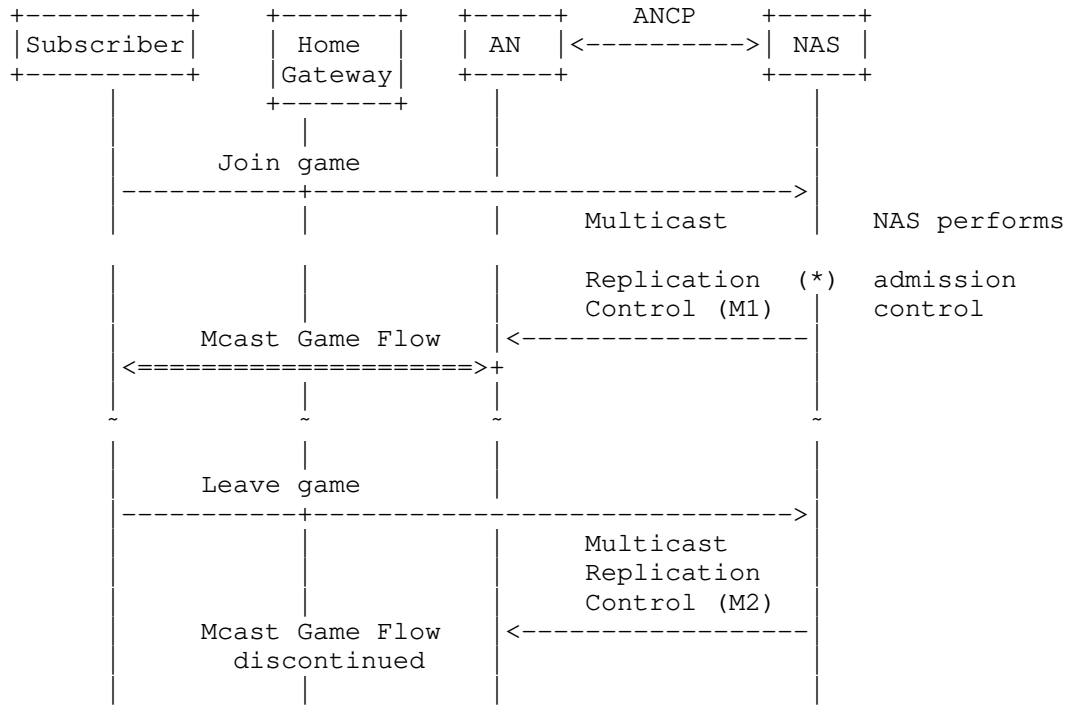

(*) The NAS may optionally seek direction from an external Authorization/Policy Server

Figure 33: NAS-Initiated Flows For On-Line Gaming

Multicast Replication Control message (M1) in Figure 34 looks like the message in Figure 26 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 = 1) with IPv4 (1) group address 233.252.0.100.

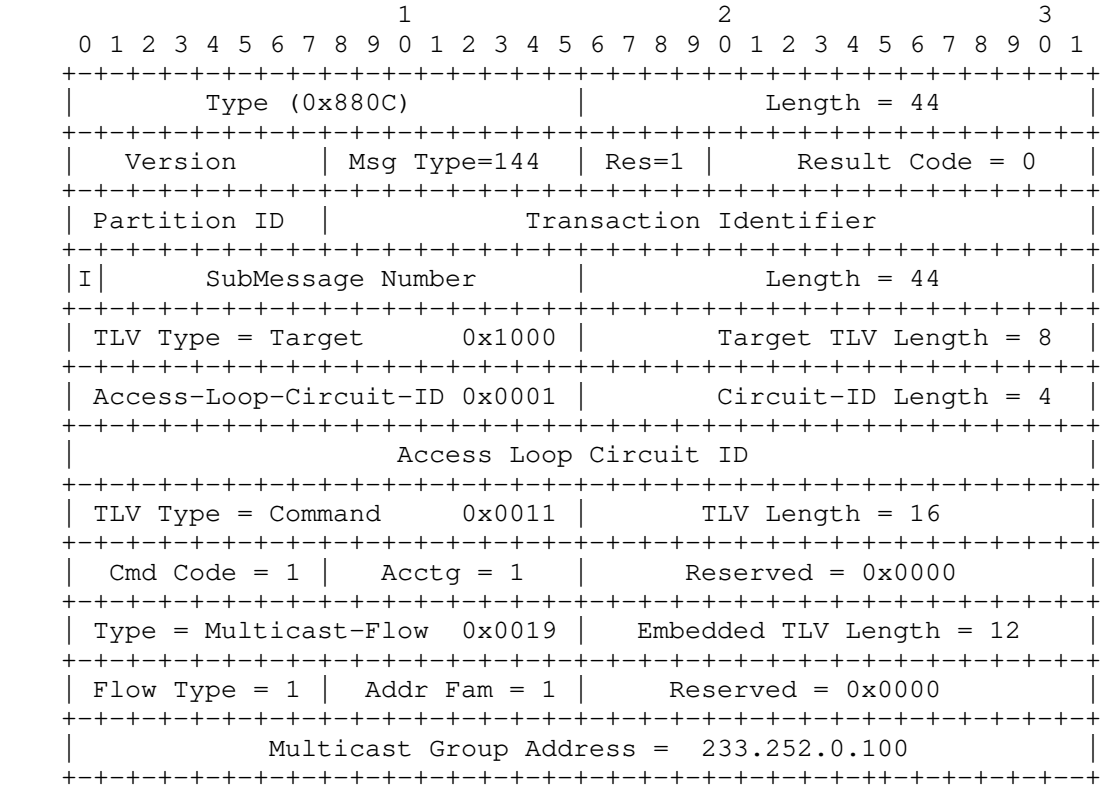

Figure 34: Enabling The Subscriber To Join An On-Line Game

Message M2 terminating the flow when the subscriber leaves the game looks the same as the message in Figure 34 with two exceptions: the Command Code becomes "Delete" (2), and Accounting is set to "No" (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 35 illustrates a message flow in a case where the NAS queries the AN about which multicast flows are active on port 10, on port 11 and on port 20 of the AN.

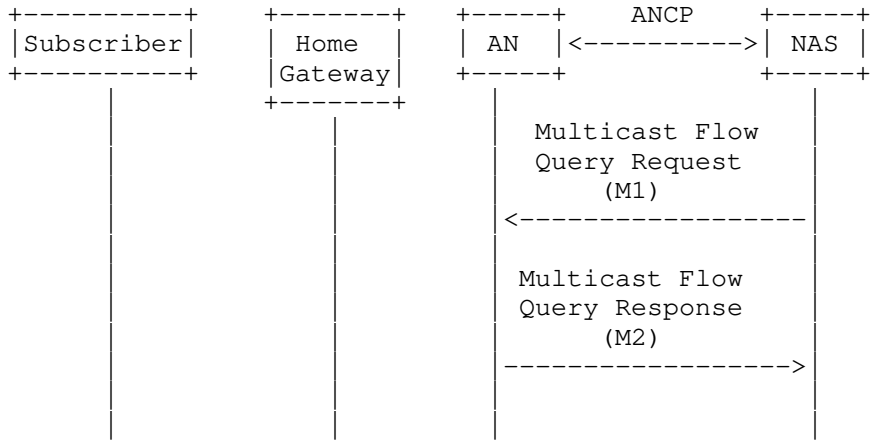

Figure 35: Per-Port Multicast Flow Reporting

The Multicast Flow Query Request message (M1) is illustrated in Figure 36. 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.

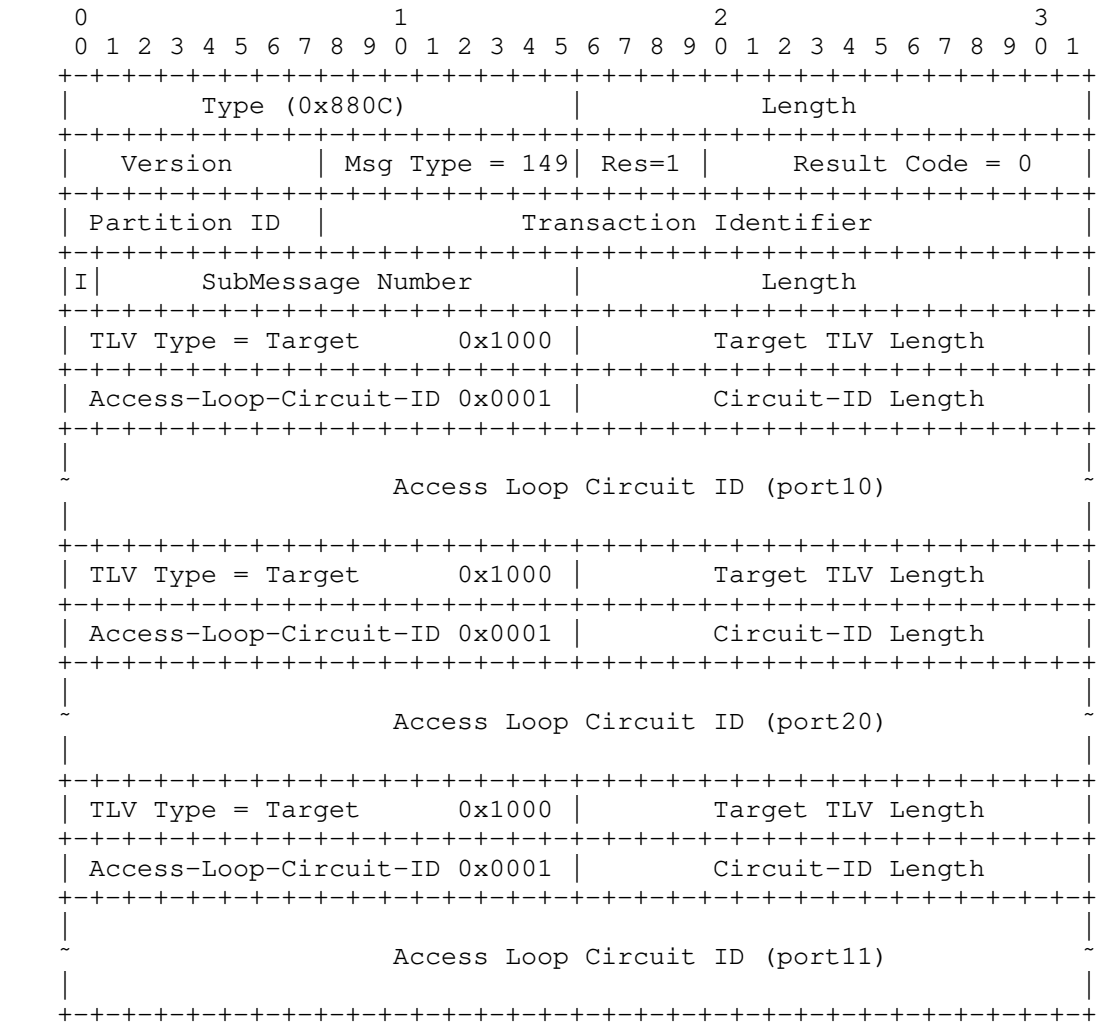

Figure 36: Multicast Flow Query Request Message For Per-Port Multicast Flow Reporting

The Multicast Flow Query Response message (M2) is illustrated in Figure 37. It indicates that there is one active multicast flow [(192.0.2.1, 233.252.0.4)] on port 10, no active multicast flow on port 20 and two active multicast flows [(192.0.2.1, 233.252.0.4) and (192.0.2.2, 233.252.0.10)] on port 11.

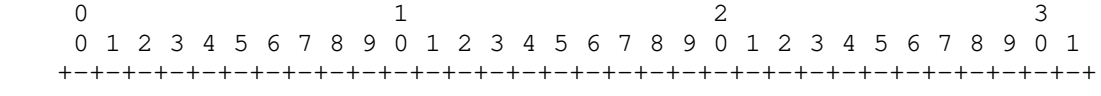

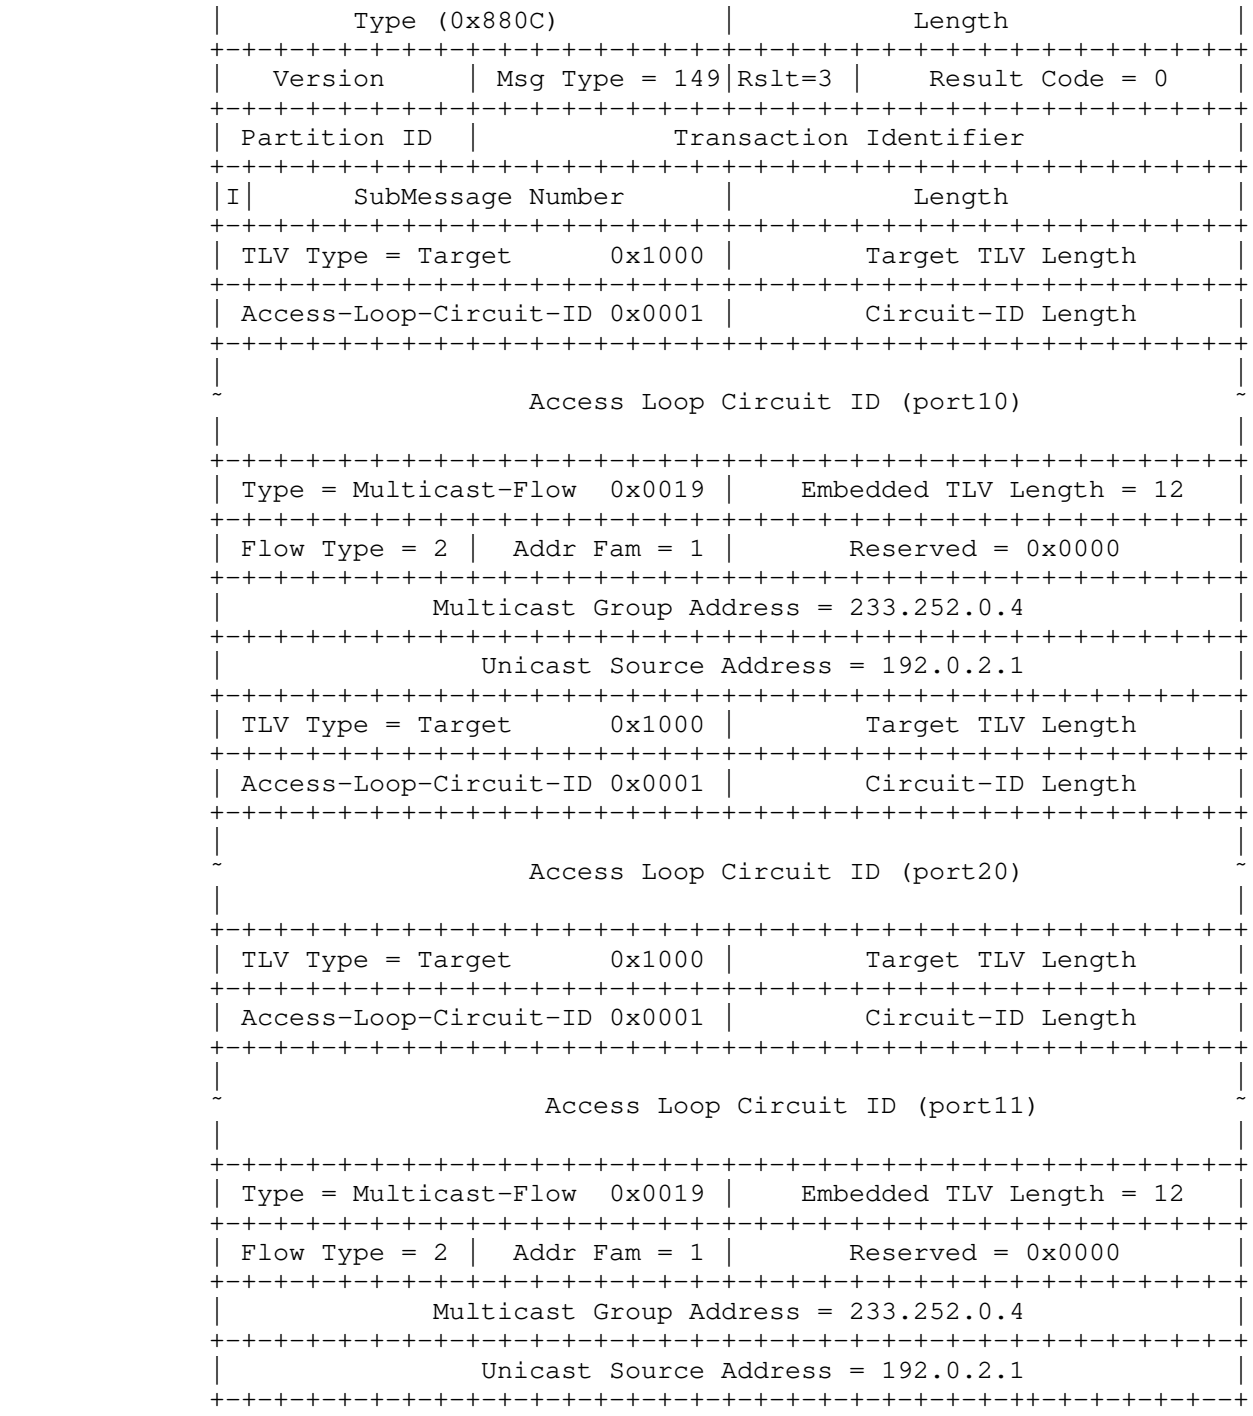
Type (0x880C)
Length
Version
Msg Type = 149
Rslt=3
Result Code = 0
Partition ID
Transaction Identifier
I
SubMessage Number
Length
TLV Type = Target 0x1000
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port10)
Type = Multicast-Flow 0x0019
Embedded TLV Length = 12
Flow Type = 2
Addr Fam = 1
Reserved = 0x0000
Multicast Group Address = 233.252.0.4
Unicast Source Address = 192.0.2.1
TLV Type = Target 0x1000
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port20)
TLV Type = Target 0x1000
Target TLV Length
Access-Loop-Circuit-ID 0x0001
Circuit-ID Length
Access Loop Circuit ID (port11)
Type = Multicast-Flow 0x0019
Embedded TLV Length = 12
Flow Type = 2
Addr Fam = 1
Reserved = 0x0000
Multicast Group Address = 233.252.0.4
Unicast Source Address = 192.0.2.1

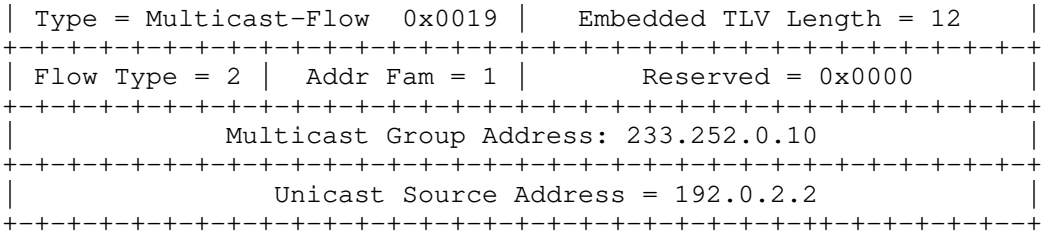

Figure 37: Multicast Flow Query Response message For Per- Port Multicast Flow Reporting

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
Cisco Systems
181 Bay Street
Toronto, ON M5J 2T3
Canada

Email: robmgl@cisco.com

Tom Taylor
Huawei Technologies
Ottawa
Canada

Email: tom.taylor.stds@gmail.com

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 28, 2011

S. Wadhwa
Alcatel-Lucent
J. Moisand
Juniper Networks
T. Haag
Deutsche Telekom
N. Voigt
Nokia Siemens Networks
T. Taylor, Ed.
Huawei Technologies

April 26, 2011

Protocol for Access Node Control Mechanism in Broadband Networks
draft-ietf-ancp-protocol-17

Abstract

This document describes the Access Node Control Protocol (ANCP). ANCP operates 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. Use cases for ANCP are documented in RFC 5851. As well as describing the base ANCP protocol, this document specifies capabilities for Digital Subscriber Line (DSL) topology discovery, line configuration, and remote line connectivity testing. The design of ANCP allows for protocol extensions in other documents if they are needed to support other use cases and other access technologies.

ANCP is based on GSMPv3 (RFC 3292), but with many modifications and extensions, to the point that the two protocols are not interoperable. For this reason, ANCP was assigned a separate version number to distinguish it.

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). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on October 28, 2011.

Table of Contents

1. Introduction

This draft defines a new protocol, the Access Node Control Protocol (ANCP), to realize a control plane between a service-oriented layer 3 edge device (the Network Access Server, NAS) and a layer 2 Access Node (e.g., Digital Subscriber Line Access Module, DSLAM) in order to perform operations related to quality of service (QoS), services, and subscriptions. The requirements for ANCP and the context within which it operates are described in [RFC5851].

ANCP provides its services to control applications operating in the AN and NAS respectively. This relationship is shown in Figure 1. Specification of the control applications is beyond the scope of this document, but informative partial descriptions are provided as necessary to give a context for the operation of the protocol.

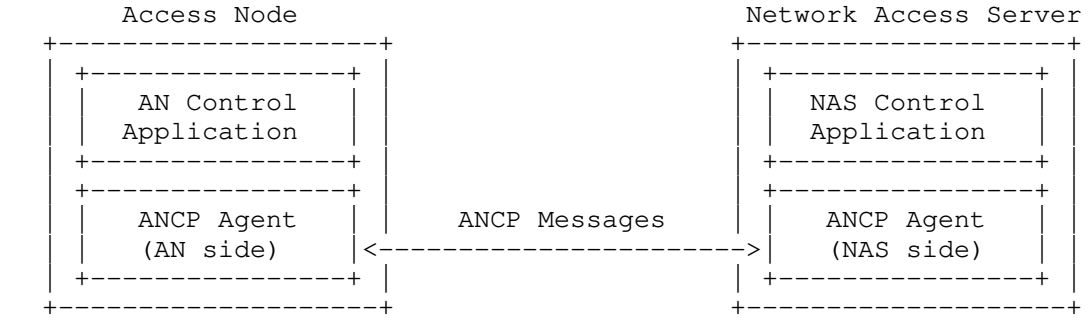

Figure 1: Architectural Context For the Access Node Control Protocol

At various points in this document, information flows between the control applications and ANCP are described. The purpose of such descriptions is to clarify the boundary between this specification and, for example, [TR-147]. There is no intention to place limits on the degree to which the control application and the protocol implementation are integrated.

This specification specifies ANCP transport over TCP/IP. TCP encapsulation for ANCP is as defined in Section 3.2.

The organization of this document is as follows:

- The next two sub-sections introduce some terminology that will be useful in understanding the rest of the document.
- Section 2 provides a description of the access networks within which ANCP will typically be deployed.

- Section 3 specifies generally applicable aspects of the ANCP protocol.
- Section 4 specifies some messages and TLVs intended for use by multiple capabilities spanning multiple technologies.
- Section 5 and the three following sections describe and specify the ANCP implementation of three capabilities applicable to the control of DSL access technology: topology discovery, line configuration, and remote line connectivity testing.
- Section 9 is the IANA Considerations section. This section defines a number of new ANCP-specific registries, as well as the joint GSMP/ANCP version registry mentioned below.
- Section 10 addresses security considerations relating to ANCP, beginning with the requirements stated in [RFC5713].

1.1. Historical Note

Initial implementations of the protocol that became ANCP were based on GSMPv3 [RFC3292]. The ANCP charter required the Working Group to develop its protocol based on these implementations. In the end, ANCP introduced so many extensions and modifications to GSMPv3 that the two protocols are not interoperable. Nevertheless, although this specification has no normative dependencies on [RFC3292], the mark of ANCP's origins can be seen in the various unused fields within the ANCP message header.

Early in ANCP's development the decision was made to use the same TCP port and encapsulation as GSMPv3, and by the time ANCP was finished it was too late to reverse that decision because of existing implementations. As a result, it is necessary to have a way for an ANCP peer to quickly distinguish ANCP from GSMP during initial adjacency negotiations. This has been provided by a joint registry of GSMP and ANCP version numbers. GSMP has version numbers 1 through 3. ANCP has the initial version number 50.

1.2. Requirements Language

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].

1.3. Terminology

This section repeats some definitions from [RFC5851], but also adds definitions for terms used only in this document.

Access Node (AN): [RFC5851] Network device, usually located at a service provider central office or street cabinet that terminates access (local) loop connections from subscribers. In case the access loop is a Digital Subscriber Line (DSL), the Access Node provides DSL signal termination, and is referred to as a DSL Access Multiplexer (DSLAM).

Network Access Server (NAS): [RFC5851] Network element which aggregates subscriber traffic from a number of Access Nodes. The NAS is an enforcement point for policy management and IP QoS in the access network. It is also referred to as a Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS).

Home Gateway (HGW): Network element that connects subscriber devices to the Access Node and the access network. In the case of DSL, the Home Gateway is a DSL network termination that may operate either 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).

ANCP agent: A logical entity that implements the ANCP protocol in the Access Node (AN-side) or NAS (NAS-side).

Access Node control adjacency: (modified from [RFC5851]) the relationship between the AN-side ANCP agent and the NAS-side ANCP agent for the purpose of exchanging Access Node Control Protocol messages. The adjacency may either be up or down, depending on the result of the Access Node Control adjacency protocol operation.

ANCP capability: A specific set of ANCP messages, message content, and procedures required to implement a specific use case or set of use cases. Some ANCP capabilities are applicable to just one access technology while others are technology independent. The capabilities applicable to a given ANCP adjacency are negotiated during adjacency startup.

Type-Length-Value (TLV): a data structure consisting of a sixteen-bit type field, a sixteen-bit length field, and a variable-length value field padded to the nearest 32-bit word boundary, as described in Section 3.6.2. The value field of a TLV can contain other TLVs. An IANA registry is maintained for values of the ANCP TLV Type field.

Net data rate: [RFC5851] defined by ITU-T G.993.2 [G.993.2], Section 3.39, i.e., the portion of the total data rate that can be used to transmit user information (e.g., ATM cells or Ethernet frames). It excludes overhead that pertains to the physical transmission mechanism (e.g., trellis coding in the case of DSL). It includes

TPS-TC (Transport Protocol Specific - Transmission Convergence) encapsulation; this is zero for ATM encapsulation, and non-zero for 64/65 encapsulation.

Line rate: [RFC5851] defined by ITU-T G.993.2. It contains the complete overhead including Reed-Solomon and trellis coding.

DSL multi-pair bonding: method for bonding (or aggregating) multiple xDSL lines into a single bi-directional logical link, henceforth referred to in this draft as "DSL bonded circuit". DSL "multi-pair" bonding allows an operator to combine the data rates on two or more copper pairs, and deliver the aggregate data rate to a single customer. ITU-T recommendations G.998.1 and G.998.2 respectively describe ATM and Ethernet based multi-pair bonding.

2. Broadband Access Aggregation

2.1. ATM-based Broadband Aggregation

The end to end DSL network consists of network service provider (NSP) and application service provider (ASP) networks, regional/access network, and customer premises network. Figure 2 shows ATM broadband access network components.

The regional/access network consists of the regional network, Network Access Server (NAS), and the access network as shown in Figure 2. Its primary function is to provide end-to-end transport between the customer premises and the NSP or ASP.

The Access Node terminates the DSL signal. It may be in the form of a DSLAM in the central office, or a remote DSLAM, or a Remote Access Multiplexer (RAM). The Access Node is the first point in the network where traffic on multiple DSL lines will be aggregated onto a single network.

The NAS performs multiple functions in the network. The NAS is the aggregation point for subscriber traffic. It provides aggregation capabilities (e.g. IP, PPP, ATM) between the Regional/Access Network and the NSP or ASP. These include traditional ATM-based offerings and newer, more native IP-based services. This includes support for Point-to-Point Protocol over ATM (PPPoA) and PPP over Ethernet (PPPoE), as well as direct IP services encapsulated over an appropriate layer 2 transport.

Beyond aggregation, the NAS is also the enforcement point for policy management and IP QoS in the regional/access networks. To allow IP QoS support over an existing non-IP-aware layer 2 access network

without using multiple layer 2 QoS classes, a mechanism based on hierarchical scheduling is used. This mechanism, defined in [TR-059], preserves IP QoS over the ATM network between the NAS and the routing gateway (RG) at the edge of the subscriber network, by carefully controlling downstream traffic in the NAS, so that significant queuing and congestion does not occur further down the ATM network. This is achieved by using a diffserv-aware hierarchical scheduler in the NAS that will account for downstream trunk bandwidths and DSL synchronization rates.

[RFC5851] provides detailed definitions of the functions of each network element in the broadband reference architecture.

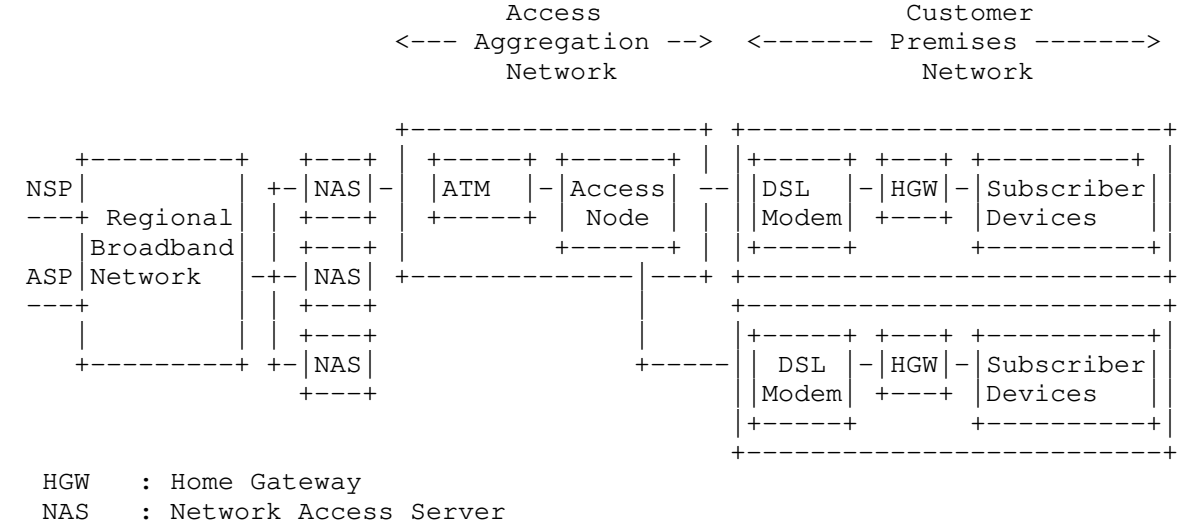

Figure 2: ATM Broadband Aggregation Topology

2.2. Ethernet-Based Broadband Aggregation

The Ethernet aggregation network architecture builds on the Ethernet bridging/switching concepts defined in IEEE 802. The Ethernet aggregation network provides traffic aggregation, class of service distinction, and customer separation and traceability. VLAN tagging defined in IEEE 802.1Q and being enhanced by IEEE 802.1ad is used as standard virtualization mechanism in the Ethernet aggregation network. The aggregation devices are "provider edge bridges" defined in IEEE 802.ad.

Stacked VLAN tags provide one possible way to create equivalent of "virtual paths" and "virtual circuits" in the aggregation network. The "outer" vlan can be used to create a form of "virtual path"

between a given DSLAM and a given NAS. "Inner" VLAN tags create a form of "virtual circuit" on a per DSL line basis. This is the 1:1 VLAN allocation model. An alternative model is to bridge sessions from multiple subscribers behind a DSLAM into a single VLAN in the aggregation network. This is the N:1 VLAN allocation model. Section 1.6 of [TR-101] provides brief definitions of these two models, while section 2.5.1 describes them in more detail.

3. Access Node Control Protocol -- General Aspects

This section specifies aspects of the Access Node Control Protocol (ANCP) that are generally applicable.

3.1. Protocol Version

ANCP messages contain an 8-bit protocol version field. For the protocol version specified in this document, the value of that field MUST be set to 50.

3.2. ANCP Transport

This document specifies the use of TCP / IPSec+IKEv2 / IP for transport of ANCP messages. For further discussion of the use of IPSec + IKEv2 see Section 10. The present section deals with the TCP aspects. Other specifications may introduce additional transports in the future.

> In the case of ATM access, a separate PVC (control channel) capable of transporting IP MAY be configured between NAS and the AN for ANCP messages.
>
> In the case of an Ethernet access/aggregation network, a typical practice is to send the Access Node Control Protocol messages over a dedicated Ethernet virtual LAN (VLAN) using a separate VLAN identifier (VLAN ID).

When transported over TCP, ANCP messages MUST use an encapsulation consisting of a four-byte header field prepended to the ANCP message as shown in Figure 3.

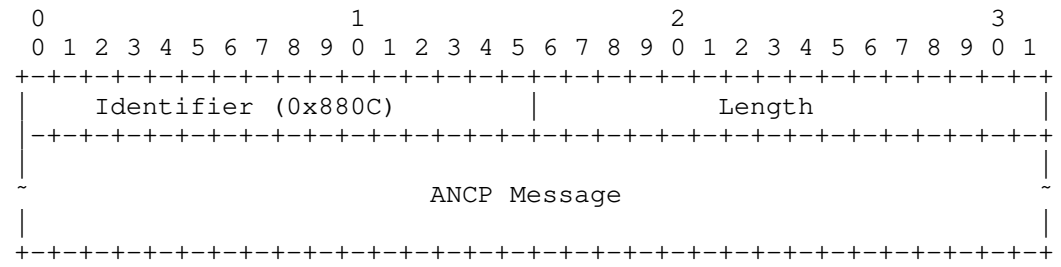

Figure 3: Encapsulation of ANCP Messages Over TCP/IP

The fields of the encapsulating header are as follows:

Identifier (16 bits): This identifies a GSMP or ANCP message. It MUST be set to 0x880C.

Length (16 bits): Total length of the ANCP message in bytes, not including the 4-byte encapsulating header.

The Access Node MUST initiate the TCP session to the NAS, using destination port 6068.

This is necessary to avoid static address provisioning on the NAS for all the ANs that are being served by the NAS. It is easier to configure a given AN with the single IP address of the NAS that serves the AN.

The NAS MUST listen on port 6068 for incoming connections from the Access Nodes.

In the event of an ANCP transport protocol failure, all pending ANCP messages destined to the disconnected recipient SHOULD be discarded until the transport connection is re-established.

3.3. Encoding of Text Fields

In ANCP, all text fields use UTF-8 encoding [RFC3629]. Note that US ASCII characters have the same representation when coded as UTF-8 as they do when coded according to [US_ASCII].

When extracting text fields from a message, the ANCP agent MUST NOT assume that the fields are zero-terminated.

3.4. Treatment of Reserved and Unused Fields

ANCP messages contain a number of fields that are unused or reserved. Some fields are always unused (typically because they were inherited

from GSMPv3 but are not useful in the ANCP context). Others are reserved in the current specification, but are provided for flexibility in future extensions to ANCP. Both reserved and unused fields MUST be set to zeroes by the sender and MUST be ignored by the receiver.

Unused bits in a flag field are shown in figures as 'x'. The above requirement (sender set to zero, receiver ignore) applies to such unused bits.

3.5. The ANCP Adjacency Protocol

ANCP uses the adjacency protocol to synchronize the NAS and Access Nodes and maintain the ANCP session. After the TCP connection is established, adjacency protocol messages MUST be exchanged as specified in this section. ANCP messages other than adjacency protocol messages MUST NOT be sent until the adjacency protocol has achieved synchronization.

3.5.1. ANCP Adjacency Message Format

The ANCP adjacency message format is shown in Figure 4 below.

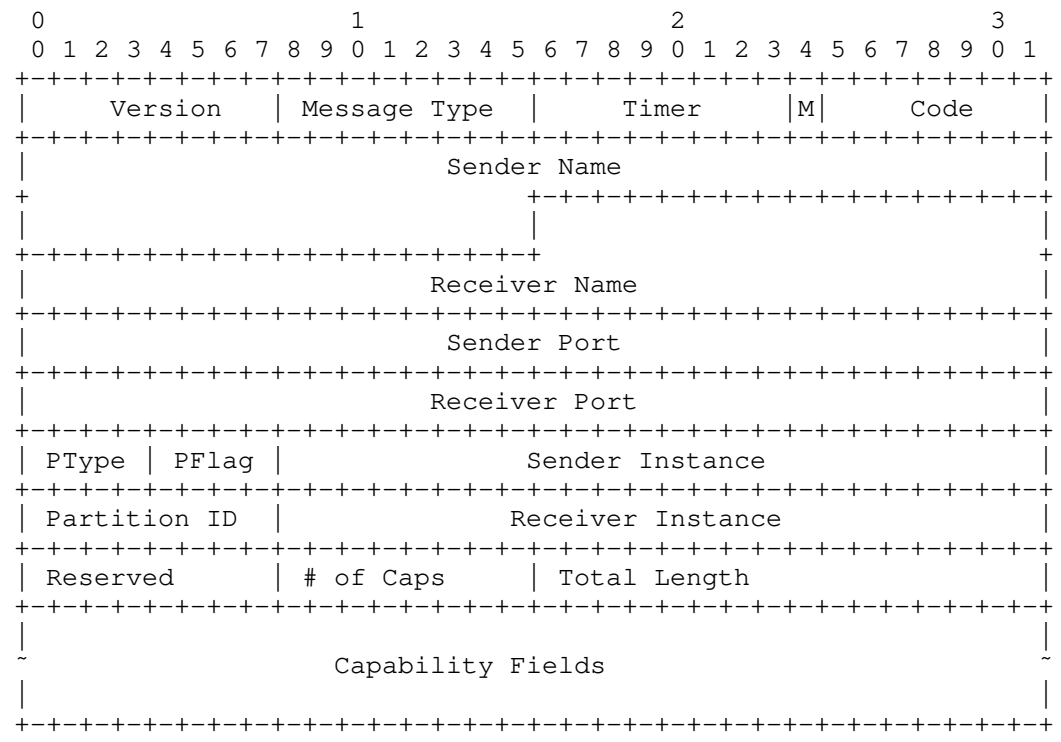

Figure 4: ANCP Adjacency Message Format

The fields of the ANCP adjacency message are as follows:

Version (8 bits): ANCP version, which is subject to negotiation. This is the key parameter by means of which ANCP messages can be distinguished from GSMP messages received over the same port.

Message Type (8 bits): always has value 10 (adjacency protocol).

Timer (8 bits): The Timer field is used to negotiate the timer value used in the adjacency protocol with the peer. The timer specifies the nominal time between periodic adjacency protocol messages. It is a constant for the duration of an ANCP session. The Timer field is specified in units of 100ms, with a default value of 250 (i.e., 25 seconds).

M-flag (one bit): used in the SYN message to prevent the NAS from synchronizing with another NAS, and the AN from synchronizing with another AN. In the SYN message, always set to 1 by the NAS, and to 0 by the AN. In other adjacency message types, always set to 0 by the sender and ignored by the receiver.

Code (7 bits): the adjacency protocol message type. It MUST have one of the following values:

Code = 1: SYN;

Code = 2: SYNACK;

Code = 3: ACK;

Code = 4: RSTACK.

Sender Name (48 bits): For the SYN, SYNACK, and ACK messages, is the identifier of the entity sending the message. The Sender Name is a 48-bit quantity that is unique within the operational context of the device. A 48-bit IEEE 802 MAC address, if available, may be used for the Sender Name. If the Ethernet encapsulation is used the Sender Name MUST be the Source Address from the MAC header. For the RSTACK message, the Sender Name field is set to the value of the Receiver Name field from the incoming message that caused the RSTACK message to be generated.

Receiver Name (48 bits) for the SYN, SYNACK, and ACK messages, is the name of the entity that the sender of the message believes is at the far end of the link. If the sender of the message does not know the name of the entity at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Name field is set to the value of the Sender Name field from the incoming message that caused the RSTACK message to be generated.

Sender Port (32 bits): For the SYN, SYNACK, and ACK messages, is the local port number of the link across which the message is being sent. For the RSTACK message, the Sender Port field is set to the value of the Receiver Port field from the incoming message that caused the RSTACK message to be generated.

Receiver Port (32 bits): For the SYN, SYNACK, and ACK messages, is what the sender believes is the local port number for the link, allocated by the entity at the far end of the link. If the sender of the message does not know the port number at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Port field is set to the value of the Sender Port field from the incoming message that caused the RSTACK message to be generated.

PType (4 bits): PType is used to specify if partitions are used and how the Partition ID is negotiated.

Type of partition being requested:

0 - no partition;

1 - fixed partition request;

2 - fixed partition assigned.

PFlag (4 bits): used to indicate the type of partition request.

1 - new adjacency;

2 - recovered adjacency.

In case of a conflict between the peers' views of the value of PFlag, the lower value is used.

Sender Instance (24 bits): For the SYN, SYNACK, and ACK messages, is the sender's instance number for the link to the peer. It is used to detect when the link comes back up after going down or when the identity of the entity at the other end of the link changes. The instance number is a 24-bit number that is guaranteed to be unique within the recent past and to change when the link or node comes back up after going down. Zero is not a valid instance number. For the RSTACK message, the Sender Instance field is set to the value of the Receiver Instance field from the incoming message that caused the RSTACK message to be generated.

Partition ID (8 bits): field used to associate the message with a specific partition of the AN. The value of this field is negotiated during the adjacency procedure. The AN makes the final decision, but will consider a request from the NAS. If the AN does not support partitions, the value of this field MUST be 0. Otherwise, it MUST be non-zero.

Receiver Instance (24 bits): For the SYN, SYNACK, and ACK messages, is what the sender believes is the current instance number for the link, allocated by the entity at the far end of the link. If the sender of the message does not know the current instance number at the far end of the link, this field SHOULD be set to zero. For the RSTACK message, the Receiver Instance field is set to the value of the Sender Instance field from the incoming message that caused the RSTACK message to be generated.

Reserved (8 bits): reserved for use by a future version of this specification.

# of Caps: indicates the number of capability fields that follow.

Total Length: indicates the total number of bytes occupied by the capability fields that follow.

Capability Fields: Each capability field indicates one ANCP capability supported by the sender of the adjacency message. Negotiation of a common set of capabilities to be supported within the ANCP session is described below. The detailed format of a capability field is shown in Figure 5 and described below.

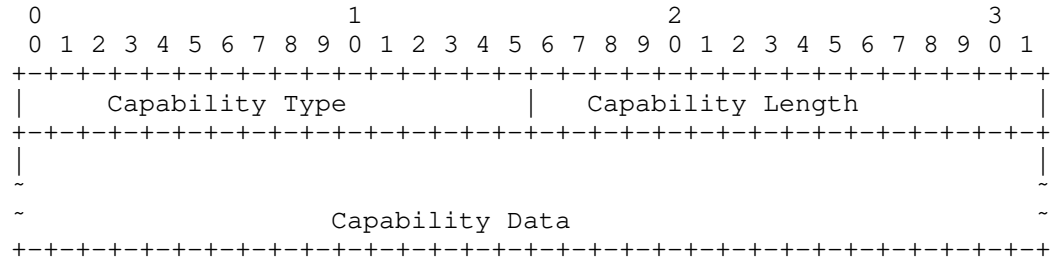

Figure 5: Capability Field

The sub-fields of this structure are as follows:

Capability Type: indicates the specific capability supported. An IANA registry exists for values of this sub-field. The values specified by this document are listed below.

Capability Length: the number of bytes of data contained in the Capability Data sub-field, excluding padding. If the definition of a particular capability includes no capability data, the value of the Capability Length sub-field is zero.

Capability Data: contains data associated with the capability as specified for that capability. If the definition of a particular capability includes no capability data, the Capability Data sub-field is absent (has zero length). Otherwise, the Capability Data sub-field MUST be padded with zeroes as required to terminate on a 4-byte word boundary. The possibility of specifying capability data provides the flexibility to advertise more than the mere presence or absence of a capability if needed.

The following capabilities are defined for ANCP as applied to DSL access:

- Capability Type : DSL Topology Discovery = 0x01

  Access technology: DSL

  Length (in bytes) : 0

  Capability Data : NULL

  For the detailed protocol specification of this capability see Section 6.

- Capability Type : DSL Line Configuration = 0x02

  Access technology: DSL

  Length (in bytes) : 0

  Capability Data : NULL

  For the detailed protocol specification of this capability see Section 7.

- Capability Type : DSL Remote Line Connectivity Testing = 0x04

  Access technology: DSL

  Length (in bytes) : 0

  Capability Data : NULL

  For the detailed protocol specification of this capability see Section 8.

In addition to the adjacency messages whose format is shown in Figure 6, ANCP adjacency procedures use the Adjacency Update message (Figure 6) to inform other NASs controlling the same AN partition when a particular NAS joins or loses an adjacency with that partition.

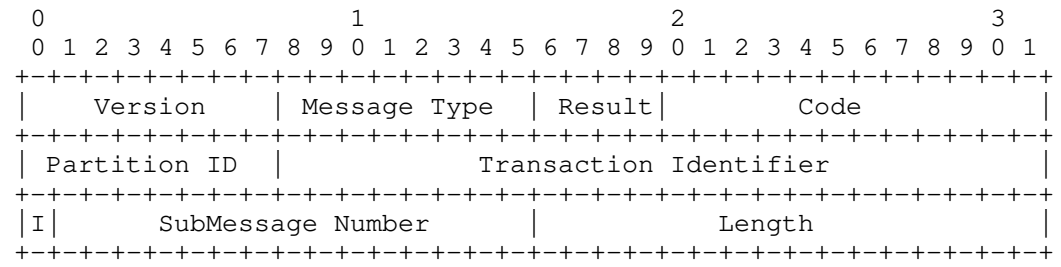

Figure 6: The Adjacency Update Message

The Adjacency Update message is identical to the general ANCP message header described in Section 3.6, but the field settings are in part specific to the Adjacency Update message. The fields in this message are as follows:

Version (8 bits): the ANCP version negotiated and running in this adjacency.

Message Type (8 bits): always 85.

Result (4 bits): set to Ignore (0).

Code (12 bits): set to the total number of adjacencies currently established on this partition, from the point of view of the AN.

Partition ID (8 bits): the partition identifier of the partition for which this notification is being sent.

Transaction Identifier (24 bits): MUST be set to 0.

I (1 bit), SubMessage number (15 bits): set as described in Section 3.6.1.7.

Length (16 bits): set as described in Section 3.6.1.8.

3.5.2. ANCP Adjacency Procedures

3.5.2.1. Overview

The ANCP adjacency protocol operates symmetrically between the NAS and the AN. In the absence of errors or race conditions, each peer sends a SYN message, receives a SYNACK message in acknowledgement, and completes the establishment of the adjacency by sending an ACK message. Through this exchange, each peer learns the values of the Name, Port, and Instance parameters identifying the other peer, and the two peers negotiate the values of the Version, Timer, PFlag, and

Partition ID parameters and the set of capabilities that the adjacency will support.

Once the adjacency has been established, its liveness is periodically tested. The peers engage in an ACK message exchange at a frequency determined by the negotiated value of the Timer field.

If an inconsistency, loss of contact, or protocol violation is detected, the detecting peer can force a restart of the synchronization process by sending an RSTACK message to the other end.

Once an adjacency has been established, if more than one NAS has established an adjacency to the same partition, then the AN sends an Adjacency Update message to each such NAS to let it know how many established adjacencies the partition currently supports. Similarly, if an adjacency is lost, the AN sends an Adjacency Update message to each of the remaining adjacent NASs to let them know about the change in status.

3.5.2.2. Adjacency Protocol State Machine

The adjacency protocol is described by the following rules and state tables. It begins with the sending of a SYN by each end as soon as the transport connection has been established. If at any point the operations A, B, C, or "Verify Adjacent State" defined below detect a mismatch, a log SHOULD be generated, identifying the fields concerned and the expected and received values for each.

The rules and state tables use the following operations:

- The "Record Adjacency State" operation is defined in Section 3.5.2.3.2.

- The "Verify Adjacency State" operation consists of verifying that the contents of the incoming SYNACK message match the adjacency state values previously recorded.

- The procedure "Reset the link" is defined as:

  1. Generate a new instance number for the link.

  2. Delete the peer verifier (set to zero the values of Sender Instance, Sender Port, and Sender Name previously stored by the Record Adjacency State operation).

  3. Send a SYN message(Section 3.5.2.3.1).

4. Enter the SYNSENT state.

- The state tables use the following Boolean terms and operators.

  A. The Sender Instance in the incoming message matches the value stored from a previous message by the "Record Adjacency State" operation.

  B. The Sender Instance, Sender Port, Sender Name and Partition ID fields in the incoming message match the values stored from a previous message by the "Record Adjacency State" operation.

  C. The Receiver Instance, Receiver Port, Receiver Name and Partition ID fields in the incoming message match the values of the Sender Instance, Sender Port, Sender Name and Partition ID currently sent in outgoing SYN, SYNACK, and ACK messages, except that the NAS always accepts the Partition ID value presented to it in a SYN or SYNACK message.

  "&&" Represents the logical AND operation.

  "||" Represents the logical OR operation.

  "!" Represents the logical negation (NOT) operation.

- A timer is required for the periodic generation of SYN, SYNACK, and ACK messages. The value of the timer is negotiated in the Timer field. The period of the timer is unspecified but a value of 25 seconds is suggested. Note that since ANCP uses a reliable transport protocol, the timer is unlikely to expire in any state other than ESTAB.

  There are two independent events: the timer expires, and a packet arrives. The processing rules for these events are:

  Timer Expires: Reset Timer

  If state = SYNSENT Send SYN

  If state = SYNRCVD Send SYNACK

  If state = ESTAB Send ACK

  Packet Arrives:

  If incoming message is an RSTACK:

If (A && C && !SYNSENT) Reset the link

Else discard the message.

If incoming message is a SYN, SYNACK, or ACK:

Response defined by the following State Tables.

If incoming message is any other ANCP message and state != ESTAB:

Discard incoming message.

If state = SYNSENT Send SYN (Note 1)

If state = SYNRCVD Send SYNACK (Note 1)

Note 1: No more than two SYN or SYNACK messages should be sent within any time period of length defined by the timer.

- State synchronisation across a link is considered to be achieved when the protocol reaches the ESTAB state. All ANCP messages, other than adjacency protocol messages, that are received before synchronisation is achieved will be discarded.

3.5.2.2.1. State Tables

State: SYNSENT

Condition	Action	New State
SYNACK && C	Update Peer Verifier; Send ACK	ESTAB
SYNACK && !C	Send RSTACK	SYNSENT
SYN	Update Peer Verifier; Send SYNACK	SYNRCVD
ACK	Send RSTACK	SYNSENT

State: SYNRCVD

Condition	Action	New State
SYNACK && C	Verify Adjacency State; Send ACK	ESTAB
SYNACK && !C	Send RSTACK	SYNRCVD
SYN	Record Adjacency State; Send SYNACK	SYNRCVD
ACK && B && C	Send ACK	ESTAB
ACK && !(B && C)	Send RSTACK	SYNRCVD

State: ESTAB

Condition	Action	New State
SYN || SYNACK	Send ACK (note 2)	ESTAB
ACK && B && C	Send ACK (note 3)	ESTAB
ACK && !(B && C)	Send RSTACK	ESTAB

Note 2: No more than two ACKs should be sent within any time period of length defined by the timer. Thus, one ACK MUST be sent every time the timer expires. In addition, one further ACK may be sent between timer expirations if the incoming message is a SYN or SYNACK. This additional ACK allows the adjacency protocol to reach synchronisation more quickly.

Note 3: No more than one ACK should be sent within any time period of length defined by the timer.

3.5.2.3. The Adjacency Protocol SYN Message

3.5.2.3.1. Action By the Sender

The SYN message is sent in accordance with the state tables just described. The sender sets the individual fields as follows:

Version: SHOULD be set to the highest version of ANCP that the sender supports.

Message Type: MUST be set to 10.

Timer: SHOULD be set to the value configured in the AN or NAS sending the message.

M Flag MUST be set to 1 by the NAS, and 0 by the AN.

Code: MUST be set to 1 (SYN).

Sender Name set as described in Section 3.5.1.

Receiver Name: SHOULD be set to 0.

Sender Port set as described in Section 3.5.1.

Receiver Port: SHOULD be set to 0.

PType: set according to the following rules:

Settings by the AN:

0 - the AN does not support partitions;

2 - the value of Partition ID contained in this message is assigned to the current partition.

Settings by the NAS:

0 - the NAS leaves the decision on partitioning to the AN (RECOMMENDED setting);

1 - the NAS requests that the AN use the value of Partition ID contained in this message for the current partition. The NAS MAY use this setting even if it has already received a SYN message from the AN, provided that the AN has indicated support for partitions. The NAS MUST be prepared to use whatever value it receives in a subsequent SYN or SYNACK message, even if this differs from the requested value.

PFlag: set to the mode of adjacency setup (new adjacency vs. recovered adjacency) requested by the sender. Warning: setting PFlag=1 runs the risk of state mismatch because ANCP does not provide the means for the NAS to audit the current state of the AN.

Sender Instance: set as described in Section 3.5.1.

Partition ID: MUST be set to 0 if PType=0, otherwise set to the assigned or requested partition identifier value.

Receiver Instance: SHOULD be set to 0.

# of Caps: MUST be set to the number of Capability fields that follow.

Total Length: MUST be set to the total number of bytes in the Capability fields that follow.

Capability Fields: one Capability field MUST be present for each ANCP capability for which the sender wishes to advertise support.

3.5.2.3.2. Action By the Receiver

Upon receiving a validly-formed SYN message, the receiver first checks the value of the Version field. If this value is not within the range of ANCP versions that the receiver supports, the message MUST be silently ignored. Similarly, the message is silently ignored if the M-flag is 0 and the receiver is an AN, or if the M-flag is 1 and the receiver is a NAS. If these checks are passed and the receiver is in ESTAB state, it returns an ACK (as indicated by the ESTAB state table in Section 3.5.2.2.1). The contents of the ACK MUST reflect the adjacency state as previously recorded by the receiver.

Otherwise, the receiver MUST record the adjacency state as follows:

Version: the supported Version value received in the SYN message. This value MUST be used for all subsequent ANCP messages sent during the life of the adjacency.

Timer: the larger of the Timer value received in the SYN message and the value with which the receiver is configured.

Sender Name: the value of the Sender Name field in the SYN message just received.

Receiver Name: the value used by the receiver in the Sender Name field of SYN, SYNACK, and ACK messages it sends in this adjacency.

Sender Port: the value of the Sender Port field in the SYN message just received.

Receiver Port: the value used by the receiver in the Sender Port field of SYN, SYNACK, and ACK messages it sends in this adjacency.

Sender Instance: the value of the Sender Instance field in the SYN message just received.

PFlag: the lesser of the value determined by local policy and the value received in the SYN message. That is, preference is given to "0 - New adjacency" if there is a conflict.

Partition ID: if the SYN receiver is the AN, this is set to 0 if the AN does not support partitions, or to the non-zero value of the partition identifier it chooses to assign otherwise. If the SYN receiver is the NAS, this is set to the value of the Partition ID field copied from the SYN.

Receiver Instance: the value used by the receiver in the Sender Instance field of SYN, SYNACK, and ACK messages it sends in this adjacency.

Capabilities: the set of ANCP capabilities that were offered in the SYN and are supported by the receiver.

3.5.2.4. The Adjacency Protocol SYNACK Message

3.5.2.4.1. Action By the Sender

The SYNACK is sent in response to a successfully received SYN message, as indicated by the state tables. The Version, Timer, PFlag, and Partition ID fields MUST be populated with the values recorded as part of adjacency state. The # of Caps, Total Length, and Capability fields MUST also be populated in accordance with the Capabilities recorded as part of adjacency state. The remaining fields of the SYNACK message MUST be populated as follows:

Message Type: MUST be 10.

M-flag: MUST be set to 0.

Code: MUST be 2 (SYNACK).

PType: MUST be 0 if the Partition ID value is 0, or 2 if the Partition ID value is non-zero.

Sender Name: MUST be set to the Receiver Name value recorded as part of adjacency state.

Receiver Name: MUST be set to the Sender Name value recorded as part of adjacency state.

Sender Port: MUST be set to the Receiver Port value recorded as part of adjacency state.

Receiver Port: MUST be set to the Sender Port value recorded as part of adjacency state.

Sender Instance: MUST be set to the Receiver Instance value recorded as part of adjacency state.

Receiver Instance: MUST be set to the Sender Instance value recorded as part of adjacency state.

If the set of capabilities recorded in the adjacency state is empty, then after sending the SYNACK the sender MUST raise an alarm to management, halt the adjacency procedure, and tear down the TCP session if it is not being used by another adjacency. The sender MAY also terminate the IPSec security association if no other adjacency is using it.

3.5.2.4.2. Action By the Receiver

As indicated by the state tables, the receiver of a SYNACK first checks that the Receiver Name, Receiver Port, and Receiver Instance values match the Sender Name, Sender Port, and Sender Instance values it sent in SYN message that is being acknowledged. The AN also checks that the PType and Partition ID match. If any of these checks fail, the receiver sends an RSTACK as described in Section 3.5.2.6.1.

The receiver next checks whether the set of capabilities provided in the SYNACK is empty. If so, the receiver MUST raise an alarm to management and halt the adjacency procedure.

Assuming that the SYNACK passes these checks, two cases arise. The first possibility is that the receiver has already recorded adjacency state. This will occur if the SYNACK is received while the receiver is in SYNRCVD state. In this case, the Version, Timer, Sender Name, Sender Port, Sender Instance, PFlag, and capability-related fields in the SYNACK MUST match those recorded as part of adjacency state. If a mismatch is detected, the receiver sends an RSTACK. This is the "Verify Adjacency State" procedure shown in the SYNRCVD state table.

If, on the other hand, the SYNACK is received while the receiver is in SYNSENT state, the receiver MUST record session state as described in Section 3.5.2.3.2.

In either case, if the receiver is the NAS, it MUST accept the Partition ID value provided in the SYNACK, updating its recorded adjacency state if necessary.

3.5.2.5. The Adjacency Protocol ACK Message

3.5.2.5.1. Actions By the Sender

As indicated by the state tables, the ACK message is sent in a number of different circumstances. The main-line usages are as a response to SYNACK, leading directly to the ESTAB state, and as a periodic test of liveness once the ESTAB state has been reached.

The sender MUST populate the ACK from recorded adjacency state, exactly as described in Section 3.5.2.4.1. The only difference is that Code MUST be set to 3 (ACK).

3.5.2.5.2. Actions By the Receiver

The required actions by the receiver are specified by the state tables. In addition to the checks B and C, the receiver SHOULD verify that the remaining contents of the ACK match the recorded adjacency state at the receiver. If that check fails the receiver MUST send an RSTACK as described in Section 3.5.2.6.1.

Once the adjacency has been established, either peer can initiate the ACK exchange that tests for liveness. To meet the restrictions on ACK frequency laid down in the notes to the state tables, it is desirable that only one such exchange occur during any one interval. Hence if a peer receives an ACK when in ESTAB state, it MUST reply to that ACK as directed by the state tables, but SHOULD NOT initiate another ACK exchange in the same interval. To meet this objective, the receiver MUST reset its timer when it receives an ACK while in ESTAB state.

> It is, of course, possible that two exchanges happen because of race conditions.

3.5.2.6. The Adjacency Protocol RSTACK Message

3.5.2.6.1. Action By the Sender

The RSTACK is sent in response to various error conditions as indicated by the state tables. In general it leads to a restart of adjacency negotiations (although this takes a few steps when the original sender of the RSTACK is in ESTAB state).

As indicated in Section 3.5.1, the Sender Name, Port, and Instance

fields in the RSTACK MUST be copied from the Receiver, Name, Port, and Instance fields in the message that caused the RSTACK to be sent. Similarly, the Receiver identifier fields in the RSTACK MUST be copied from the corresponding Sender identifier fields in the message that triggered the RSTACK.

If the sender has recorded adjacency state, the Version, Timer, PType, PFlag, Partition ID, and capability-related fields SHOULD be set based on the recorded adjacency state. Otherwise they SHOULD be the same as the sender would send in a SYN message. The Message Type MUST be 10, the M-flag MUST be 0, and Code MUST be 4 (RSTACK).

3.5.2.6.2. Action By the Receiver

The receiver of an RSTACK MAY attempt to diagnose the problem which caused the RSTACK to be generated by comparing its own adjacency state with the contents of the RSTACK. However, the primary purpose of the RSTACK is to trigger action as prescribed by Section 3.5.2.2.

3.5.2.7. Loss of Synchronization

Loss of synchronisation MAY be declared if after synchronisation is achieved:

- no valid ANCP messages are received in any period of time in excess of three times the value of the Timer field negotiated in the adjacency protocol messages, or
- a mismatch in adjacency state is detected.

In either case the peer detecting the condition MUST send an RSTACK to the other peer as directed in Section 3.5.2.6.1, in order to initiate resynchronization.

While re-establishing synchronisation with a controller, a switch SHOULD maintain its connection state, deferring the decision about resetting the state until after synchronisation is re-established. Once synchronisation is re-established the decision about resetting the connection state SHOULD be made based on the negotiated value of PFlag.

3.6. ANCP General Message Formats

This section describes the general format of ANCP messages other than the adjacency messages. See Figure 7

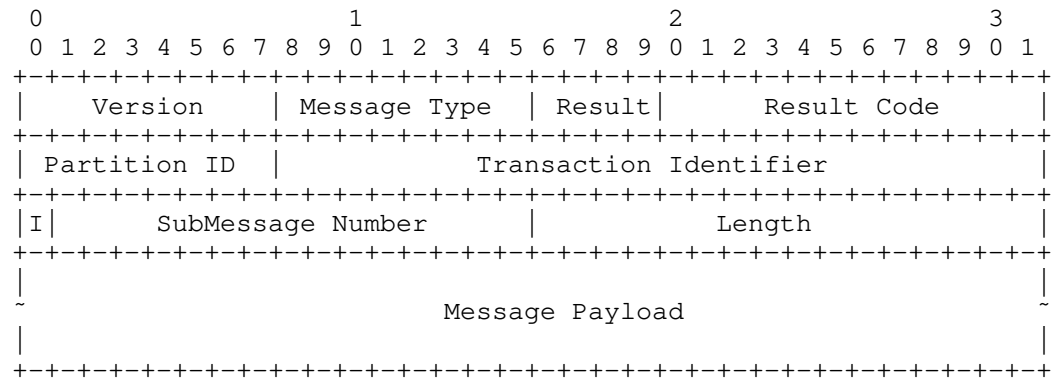

Figure 7: ANCP General Message Format

3.6.1. The ANCP Message Header

A complete explanation of the ANCP general message header fields follows.

3.6.1.1. Version Field (8 bits)

This field carries the version of the ANCP protocol that was agreed for the session during adjacency negotiation.

3.6.1.2. Message Type Field (8 bits)

This field indicates the ANCP message type. Message type values are registered in an IANA registry.

3.6.1.3. Result Field (4 bits)

In request messages, the Result field indicates the circumstances under which a response is required. ANCP specifies what Result value each request message type should have. In responses the Result field indicates either Success (0x3) or Failure (0x4) as the case may be.

Ignore: Res = 0x0 - Treat this field as a "no operation" and follow the response procedures specified for the received message type.

Nack: Res = 0x1 - Result value indicating that a response is expected to the request only in cases of failure caused during the processing of the message contents or of the contained directive(s).

AckAll: Res = 0x2 - Result value indicating that a response to the message is requested in all cases.

Success: Res = 0x3 - Result value indicating that this is a response and that the request was executed successfully. The Result Code field for a successful result is typically 0, but MAY take on other values as specified for particular message types.

Failure: Res = 0x4 - Result value indicating that this is a response and that the request was not executed successfully. The receiver of the response SHOULD take further action as indicated by the Result Code value and any diagnostic data contained in a Status-Info TLV included in the response.

3.6.1.4. Result Code Field (12 bits)

This field gives further information concerning the result in a response message. It is mostly used to pass an error code in a failure response but can also be used to give further information in a success response message or an event message. In a request message, the Result Code field is not used and MUST be set to 0x0 (No result).

A number of Result Code values are specified below. Specification of additional Result Code values in extensions or updates to this document MUST include the following information:

- Result Code value;
- One-line description;
- Where condition detected: (control application or ANCP agent);
- Further description (if any);
- Required additional information in the response message;
- Target (control application or ANCP agent at the peer that sent the original request);
- Action RECOMMENDED for the receiving ANCP agent

In addition to any suggested action in the text which follows, a count of the number of times a given non-zero Result Code value was received SHOULD be provided for management. Where an action includes resending of a request, a given request SHOULD NOT be re-sent more than once.

This document specifies the following Result Code values.

Result Code value: 0x2

* One-line description: Invalid request message
* Where condition detected: ANCP agent
* Further description: The request was a properly formed message which violates the protocol through its timing or direction of transmission. The most likely reason for this outcome in the field will be a race condition.
* Required additional information in the response message: none, if the response message is of the same type as the request. As specified in Section 4.2 if the response message is a Generic Response message.
* Target: ANCP agent at the peer that sent the original request
* Action RECOMMENDED for the receiving ANCP agent: The original request MAY be re-sent once only after a short delay. Inform the control application with appropriate identification of the failed transaction if the second attempt fails or no second attempt is made.

Result Code value: 0x6

* One-line description: One or more of the specified ports are down
* Where condition detected: control application
* Further description (if any): This Result Code value indicates a state mismatch between the NAS and AN control applications, possibly due to a race condition.
* Required additional information in the response message: if the request identified multiple access lines or the response is a Generic Response message, then the response MUST contain a Status-Info TLV encapsulating TLV(s) containing the line identifier(s) of the access lines that are not operational.
* Target: control application at the peer that sent the original request
* Action RECOMMENDED for the receiving ANCP agent: indicate the error and forward the line identifier(s) to the control

application.

Result Code value: 0x13

* One-line description: Out of resources

* Where condition detected: ANCP protocol layer or control application

* Further description: (e.g., memory exhausted, etc.). This Result Code value MUST be reported only by the AN, and indicates a condition that is probably unrelated to specific access lines (although it may be related to the specific request).

* Required additional information in the response message: none, if the response message is of the same type as the request. As specified in Section 4.2 if the response message is a Generic Response message.

* Target: ANCP agent at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: If the NAS receives this Result Code value from multiple requests for the same AN in a short interval, it SHOULD reduce the rate at which it sends requests in proportion to the rate at which requests are failing with Result Code = 19. It MAY retry individual requests. If only a specific request is failing with Result Code = 19, the ANCP agent in the NAS MAY request the control application to decompose the request into simpler components if this is possible.

Result Code value: 0x51

* One-line description: Request message type not implemented

* Where condition detected: ANCP agent

* Further description: This could indicate a mismatch in protocol version or capability state. It is also possible that support of a specific message is optional within some ANCP capability.

* Required additional information in the response message: none, if the response message is of the same type as the request. As specified in Section 4.2 if the response message is a Generic Response message.

* Target: ANCP agent at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: If the receiver of this Result Code value expects that support of the message type concerned is mandatory according to the capabilities negotiated for the session, it MAY re-send the message in case the message was corrupted in transit the first time. If that fails, and use of the message type cannot be avoided, the ANCP agent MAY reset the adjacency by sending an RSTACK adjacency message (Section 3.5.2.6.1) where PType is set to 0 and Sender and Receiver Name, Port, and Instance are taken from recorded adjacency state. If a reset does not eliminate the problem, the receiving ANCP agent SHOULD raise an alarm to management and then cease to operate.

Result Code value: 0x53

* One-line description: Malformed message

* Where condition detected: ANCP agent

* Further description: This could be the result of corruption in transit, or an error in implementation at one end or the other.

* Required additional information in the response message: none, if the response message is of the same type as the request. As specified in Section 4.2 if the response message is a Generic Response message.

* Target: ANCP agent at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: The request SHOULD be re-sent once to eliminate the possibility of in-transit corruption.

Result Code value: 0x54

* One-line description: Mandatory TLV missing

* Where condition detected: ANCP agent

* Further description: none.

* Required additional information in the response message: the response message MUST contain a Status-Info message that encapsulates an instance of each missing mandatory TLV, where the length is set to zero and the value field is empty (i.e., only the four-byte TLV header is present).

* Target: ANCP agent at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: resend the message with the missing TLV(s), if possible. Otherwise, report the error to the control application with an indication of the missing information required to construct the missing TLV(s).

Result Code value: 0x55

* One-line description: Invalid TLV contents

* Where condition detected: ANCP agent

* Further description: the contents of one or more TLVs in the request do not match the specifications provided for the those TLVs.

* Required additional information in the response message: the response MUST contain a Status-Info TLV encapsulating the erroneous TLVs copied from the original request.

* Target: ANCP agent at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: correct the error and resend the request, if possible. Otherwise, report the error to the control application with an indication of the erroneous information associated with the invalid TLV(s).

Result Code value: 0x500

* One-line description: One or more of the specified ports do not exist

* Where condition detected: control application

* Further description (if any): this may indicate a configuration mismatch between the AN and the NAS or AAA.

* Required additional information in the response message: if the request identified multiple access lines or the response is a Generic Response message, then the response MUST contain a Status-Info TLV encapsulating TLV(s) containing the rejected line identifier(s).

* Target: control application at the peer that sent the original request

* Action RECOMMENDED for the receiving ANCP agent: indicate the error and forward the line identifiers to the control application.

3.6.1.5. Partition ID (8 bits)

The Partition ID field MUST contain the value that was negotiated for Partition ID during the adjacency procedure as described above.

3.6.1.6. Transaction ID (24 bits)

The Transaction ID is set by the sender of a request message to associate a response message with the original request message. Unless otherwise specified for a given message type, the Transaction ID in request messages MUST be set to a value in the range (1, 2^24 - 1). When used in this manner, the Transaction ID sequencing MUST be maintained independently for each message type within each ANCP adjacency. Furthermore, it SHOULD be incremented by 1 for each new message of the given type, cycling back to 1 after running the full range. For event messages, the Transaction ID SHOULD be set to zero.

Unless otherwise specified, the default behaviour for all ANCP responses is that the value of the Transaction ID MUST be copied from the corresponding request message.

3.6.1.7. I flag and SubMessage Number (1 + 15 bits)

In GSMPv3 these provide a mechanism for message fragmentation. Because ANCP uses TCP transport, this mechanism is unnecessary. An ANCP agent MUST set the I Flag and subMessage Number fields to 1 to signify "no fragmentation".

3.6.1.8. Length (16 bits)

This field MUST be set to the length of the ANCP message in bytes, including its header fields and message body but excluding the four-byte encapsulating header defined in Section 3.2.

3.6.2. The ANCP Message Body

The detailed contents of the message payload portion of a given ANCP message can vary with the capability in the context of which it is being used. However, the general format consists of zero or more fixed fields, followed by a variable amount of data in the form of Type-Length-Value (TLV) data structures.

The general format of a TLV is shown in Figure 8:

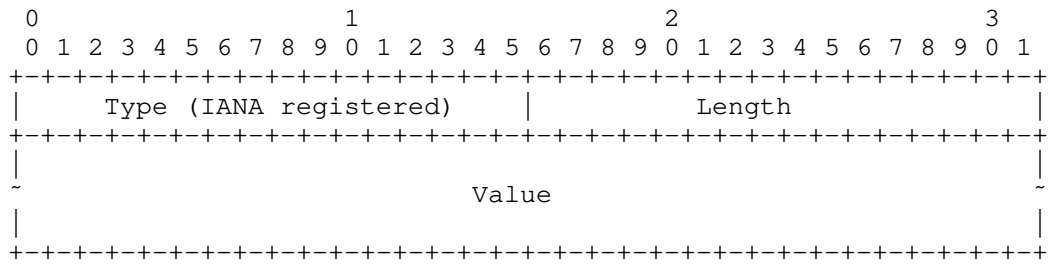

Figure 8: General TLV Format

The fields of a TLV are defined as follows:

Type (16 bits): The TLV Type is an unsigned value identifying the TLV type and nature of its contents. An IANA registry has been established for ANCP TLV Type codes.

Length (16 bits): The number of bytes of data in the Value field of the TLV, excluding any padding required to bring this TLV to a 4-byte word boundary (see "Value" below). If a TLV contains other TLVs, any padding in the contained TLVs MUST be included in the value of Length. Depending on the specification of the TLV, the value of Length can be zero, a constant for all instances of the TLV, or a varying quantity.

Value (variable): The actual data carried by the TLV, if any. The value field in each TLV MUST be padded with zeroes as required to align with a 4-byte word boundary. The Value field of a TLV MAY include fixed fields and/or other TLVs.

Unless otherwise specified, TLVs MAY be added to a message in any order. If the recipient of a message does not understand a particular TLV, it MUST silently ignore it.

A number of TLVs are specified in the remainder of this document.

3.7. General Principles for the Design of ANCP Messages

ANCP allows for two messaging constructs to support request/response interaction:

a. The same message type is used for both the request message and the response message. The Result and Result Code field settings are used to differentiate between request and response messages.

b. The request and response messages use two different message types.

The first approach is illustrated by the protocol specifications in Section 8.4, the second by specifications in Section 6.4. The purpose of this section is to provide more details about the second approach in order to allow the use of this messaging construct for the development of additional ANCP extensions.

As Section 3.6 indicated, all ANCP messages other than adjacency messages share a common header format. When the response message type is different from that of the request, the specification of the request message will typically indicate that the Result field is set to Ignore (0x0) and provide procedures indicating explicitly when the receiver should generate a response and what message type it should use.

The Transaction ID field is used to distinguish between multiple request messages of the same type and to associate a response message to a request. Specifications of ANCP messages for applications not requiring response correlation SHOULD indicate that the Transaction ID MUST be set to zero in requests. Applications that require response correlation SHOULD refer to the Transaction ID behaviour described in Section 3.6.1.

The specification for a response message SHOULD indicate in all cases that value of the Transaction Identifier MUST be set to that of the corresponding request message. This allows the requester to establish whether or not correlation is needed (by setting a non-zero or zero value for the Transaction ID).

4. Generally Useful ANCP Messages and TLVs

This section defines two messages and a number of TLVs that could be useful in multiple capabilities. In some cases the content is under-specified, with the intention that particular capabilities spell out the remaining details.

4.1. Provisioning Message

The Provisioning message is sent by the NAS to the AN to provision information of global scope (i.e., not associated with specific access lines) on the AN. The Provisioning message has the format shown in Figure 9. Support of the Provisioning message is OPTIONAL unless the ANCP agent claims support for a capability that requires its use.

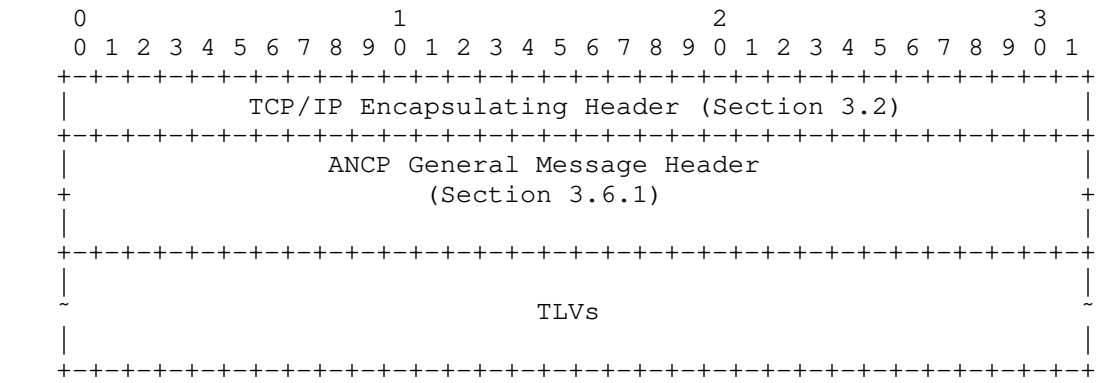

Figure 9: Format of the Provisioning Message

The message header field settings given below are REQUIRED in the Provisioning message. The remaining message header fields MUST be set as specified in Section 3.6.1. Which TLVs to carry in the Provisioning message is specified as part of the specification of the capabilities that use that message. The Provisioning message MAY be used to carry data relating to more than one capability at once, assuming that the capabilities concerned can co-exist and have all been negotiated during adjacency establishment.

Message Type: MUST be set to 93.

Result: MUST be set to 0x0 (Ignore).

Result Code: MUST be set to zero.

Transaction ID: MUST be populated with a non-zero value chosen in the manner described in Section 3.6.1.6.

If the AN can process the message successfully and accept all the provisioning directives contained in it, the AN MUST NOT send any response.

Unless otherwise specified for a particular capability, if the AN fails to process the message successfully it MUST send a Generic Response message (Section 4.2) indicating failure and providing appropriate diagnostic information.

4.2. Generic Response Message

This section defines the Generic Response message. The Generic Response message MAY be specified as the appropriate response to a message defined in an extension to ANCP, instead of a more specific

response message. As a general guideline, specification of the Generic Response message as a response is appropriate where no data needs to be returned to the peer other than a result (success or failure), plus, in the case of a failure, a code indicating the reason for failure and a limited amount of diagnostic data. Depending on the particular use case, the Generic Response message MAY be sent by either the NAS or the AN.

Support of the Generic Response message, both as sender and as receiver, is REQUIRED for all ANCP agents, regardless of what capabilities they support.

The AN or NAS MAY send a Generic Response message indicating a failure condition independently of a specific request before closing the adjacency as a consequence of that failure condition. In this case, the sender MUST set the Transaction ID field in the header and the Message Type field within the Status-Info TLV to zeroes. The receiver MAY record the information contained in the Status-Info TLV for management use.

The format of the Generic Response message is shown in Figure 10

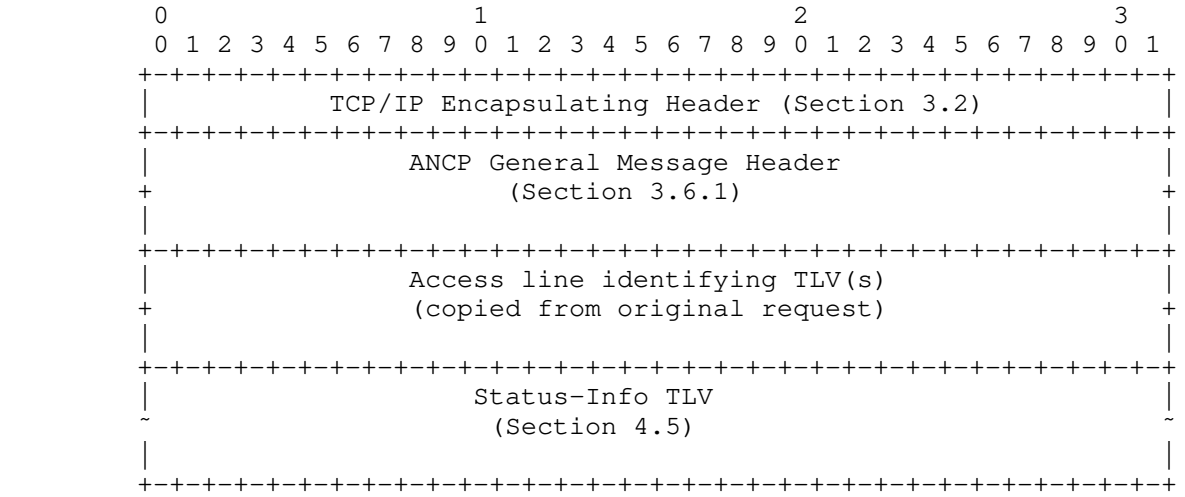

NOTE: TLVs MAY be in a different order from what is shown in this figure.

Figure 10: Structure of the Generic Response Message

This document specifies the following header fields. The remaining fields in the ANCP general message header MUST be set as specified in Section 3.6.1.

Message Type: MUST be set to 91.

Result: MUST be set to 0x3 (Success) or 0x4 (Failure).

Result Code: MUST be set to zero for success or an appropriate non-zero value for failure.

Transaction ID: MUST be copied from the message to which this message is a response.

If the original request applied to a specific access line or set of lines, the TLVs identifying the line(s) and possibly the user MUST be copied into the Generic Response message at the top level.

The Status-Info TLV MAY be present in a success response, to provide a warning as defined for a specific request message type. It MUST be present in a failure response. See Section 4.5 for a detailed description of the Status-Info TLV. The actual contents will depend on the request message type this message is responding to and the value of the Result Code field.

To prevent an infinite loop of error responses, if the Generic Response message is itself in error, the receiver MUST NOT generate an error response in return.

4.3. Target TLV

Type: 0x1000 to 0x1020 depending on the specific content. Only 0x1000 has been assigned in this specification (see below). Support of any specific variant of the Target TLV is OPTIONAL unless the ANCP agent claims support for a capability that requires its use.

Description: The Target TLV (0x1000 - 0x1020) is intended to be a general means to represent different types of objects.

Length: Variable, depending on the specific object type.

Value: Target information as defined for each object type. The Value field MAY consist of sub-TLVs.

TLV Type 0x1000 is assigned to a variant of the Target TLV representing a single access line and encapsulating one or more sub-TLVs identifying the target. Figure 11 is an example illustrating the TLV format for a single port identified by an Access-Loop-Circuit-ID TLV (0x0001) (Section 5.1.2.1).

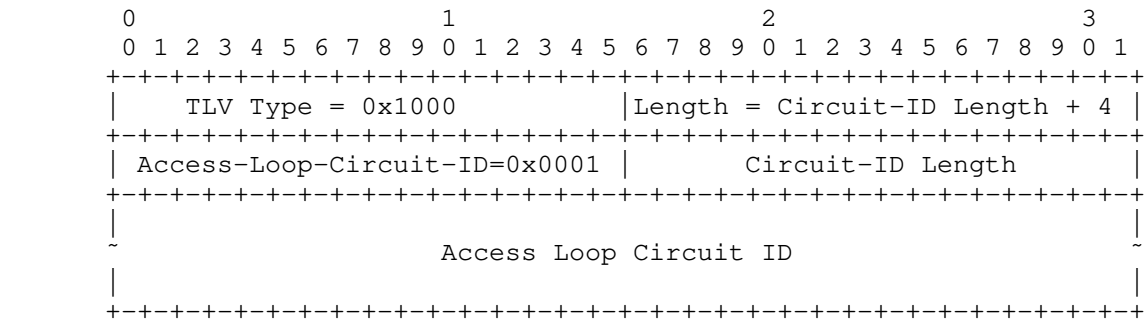

Figure 11: Example of Target TLV For Single Access Line

4.4. Command TLV

Type: 0x0011

Description: The Command TLV (0x0011) is intended to be a general means of encapsulating one or more command directives in a TLV oriented message. The semantics of the command can be specified for each message type using it. I.e., the specification of each message type that can carry the Command TLV is expected to define the meaning of the content of the payload, although re-use of specifications is, of course, permissible when appropriate. Support of any specific variant of the Command TLV is OPTIONAL unless the ANCP agent claims support for a capability that requires its use.

Length: Variable, depending on the specific contents.

Value: Command information as defined for each message type. The field MAY include sub-TLVs. The contents of this TLV MUST be specified as one "command" or alternatively a sequence of one or more "commands", each beginning with a one-byte Command Code and possibly including other data following the Command Code. An IANA registry has been established for Command Code values. This document reserves the Command Code value 0 as an initial entry in the registry.

4.5. Status-Info TLV

Name: Status-Info

Type: 0x0106

Description: The Status-Info-TLV is intended to be a general container for warning or error diagnostics relating to commands and/or requests. It is a supplement to the Result Code field in the ANCP general header. The specifications for individual message types MAY indicate the use of this TLV as part of responses, particularly for failures. As mentioned above, the Generic Response message will usually include an instance of the Status-Info TLV. Support of the Status-Info TLV, both as sender and as receiver, is REQUIRED for all ANCP agents, regardless of what capabilities they support.

Length: Variable, depending on the specific contents.

Value: The following fixed fields. In addition, sub-TLVs MAY be appended to provide further diagnostic information.

Reserved (8 bits): see Section 3.4 for handling of reserved fields.

Msg Type (8 bits): Message Type of the request for which this TLV is providing diagnostics.

Result Message Length (16 bits): Number of bytes in the error message, excluding padding, but including the language tag and delimiter. This MAY be zero if no error message is provided.

Result Message: Human-readable string providing information about the warning or error condition. The initial characters of the string MUST be a language tag as described in [RFC5646], terminated by a colon (":"). The actual text string follows the delimiter. The field is padded at the end with zeroes as necessary to extend it to a four-byte word boundary.

Section 3.6.1.4 provides recommendations for what TLVs to add in the Status-Info TLV for particular values of the message header Result Code field.

Figure 12 illustrates the Status-Info TLV.

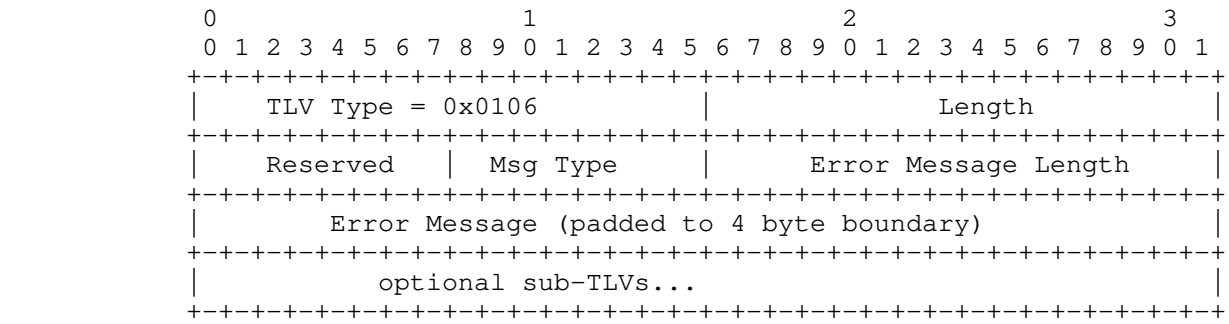

Figure 12: The Status-Info TLV

5. Introduction To ANCP Capabilities For Digital Subscriber Lines (DSL)

DSL is a widely deployed access technology for Broadband Access for Next Generation Networks. Specifications such as [TR-059], [TR-058], and [TR-092] describe possible architectures for these access networks. The scope of these specifications includes the delivery of voice, video, and data services.

The next three sections of this document specify basic ANCP capabilities for use specifically in controlling Access Nodes serving DSL access (Tech Type = 0x05). The same ANs could be serving other access technologies (e.g. Metro-Ethernet, Passive Optical Networking, WiMax), in which case the AN will also have to support the corresponding other-technology-specific capabilities. Those additional capabilities are outside the scope of the present document.

5.1. DSL Access Line Identification

Most ANCP messages involve actions relating to a specific access line. Thus it is necessary to describe how access lines are identified within those messages. This section defines four TLVs for that purpose and provides an informative description of how they are used.

5.1.1. Control Context (Informative)

Three types of identification are described in [TR-101] and provided for in the TLVs defined in this section:

- identification of an access line by its logical appearance on the user side of the Access Node;

- identification of an access line by its logical appearance on the NAS side of the Access Node; and

- identification down to the user or host level as a supplement to access line identification in one of the other two forms.

All of these identifiers originate with the AN control application, during the process of DSL topology discovery. The control application chooses which identifiers to use and the values to place into them on a line-by-line basis, based on AN configuration and deployment considerations.

Aside from its use in ANCP signalling, access line identification is also used in DHCP ([RFC2131], [RFC3315]) transactions involving hosts served by DSL. Either the AN or the NAS can serve as a DHCP relay node. [TR-101] requires the AN or NAS in this role to add access line identification in Option 82 (Information) ([RFC3046], with its IPv6 equivalent in [RFC4649]) to each DHCP request it forwards to the DHCP server. It is desirable for efficiency that the identification used in this signalling should be the same as the identification used in ANCP messages.

From the point of view of ANCP itself, the identifiers are opaque. From the point of view of the AN control application, the syntax for the user-side access line identifier is the same as specified in Section 3.9.3 of [TR-101] for DHCP Option 82. The syntax for the ASCII form of the NAS-side access line identifier will be similar.

Access line identification by logical appearance on the user side of the Access Node will always identify a DSL loop uniquely. Identification by the logical appearance on the NAS side of the Access Node is unique only if there is a one-to-one mapping between the appearances on the two sides and no identity-modifying aggregation between the AN and the NAS. In other cases, and in particular in the case of Ethernet aggregation using the N:1 VLAN model, the user-side access line identification is necessary, but the NAS-side identification is potentially useful information allowing the NAS to build up a picture of the aggregation network topology.

Additional identification down to the user or host level is intended to supplement rather than replace either of the other two forms of identification.

> Sections 3.8 and 3.9 of [TR-101] are contradictory on this point. It is assumed here that Section 3.9 is meant to be authoritative.

The user-level identification takes the form of an administered string which again is opaque at the ANCP level.

The NAS control application will use the identifying information it receives from the AN directly for some purposes. For examples, see the introductory part of Section 3.9 of [TR-101]. For other purposes, the NAS will build a mapping between the unique access line identification provided by the AN, the additional identification of the user or host (where provided), and the IP interface on a particular host. For access lines with static IP address assignment that mapping could be configured instead.

5.1.2. TLVs For DSL Access Line Identification

This section provides a normative specification of the TLVs that ANCP provides to carry the types of identification just described. The Access-Loop-Circuit-ID TLV identifies an access line by its logical appearance on the user side of the Access Node. Two alternatives, the Access-Aggregation-Circuit-ID-ASCII TLV and the Access-Aggregation-Circuit-ID-Binary TLV, identify an access line by its logical appearance on the NAS side of the Access Node. It is unlikely that a given AN uses both of these TLVs, either for the same line or for different lines, since they carry equivalent information. Finally, the Access-Loop-Remote-Id TLV contains an operator-configured string that uniquely identifies the user on the associated access line, as described in Sections 3.9.1 and 3.9.2 of [TR-101].

ANCP agents conforming to this section MUST satisfy the following requirements:

- ANCP agents MUST be able to build and send the Access-Loop-Circuit-ID TLV, the Access-Loop-Remote-Id TLV, and either the Access-Aggregation-Circuit-ID-ASCII TLV or the Access-Aggregation-Circuit-ID-Binary TLV (implementation choice), when passed the associated information from the AN control application.

- ANCP agents MUST be able to receive all four TLV types, extract the relevant information, and pass it to the control application.

- If the Access-Loop-Remote-Id TLV is present in a message, it MUST be accompanied by an Access-Loop-Circuit-ID TLV and/or an Access-Aggregation-Circuit-ID-xxx TLV with two VLAN identifiers.

  The Access-Loop-Remote-Id TLV is not enough to identify an access line uniquely on its own. As indicated above, an Access-Aggregation-Circuit-ID-xxx TLV with two VLAN identifiers may or may not identify an access line uniquely, but this is up to the control application to decide.

- If the Access-Aggregation-Circuit-ID-xxx TLV is present in a message with just one VLAN identifier, it MUST be accompanied by an Access-Loop-Circuit-ID TLV.

5.1.2.1. Access-Loop-Circuit-ID TLV

Type: 0x0001

Description: a locally administered human-readable string generated by or configured on the Access Node, identifying the corresponding access loop logical port on the user side of the Access Node.

Length: up to 63 bytes

Value: ASCII string

5.1.2.2. Access-Loop-Remote-Id TLV

Type: 0x0002

Description: an operator-configured string that uniquely identifies the user on the associated access line, as described in Sections 3.9.1 and 3.9.2 of [TR-101].

Length: up to 63 bytes

Value: ASCII string

5.1.2.3. Access-Aggregation-Circuit-ID-Binary TLV

Type: 0x0006

Description: This TLV identifies or partially identifies a specific access line by means of its logical circuit identifier on the NAS side of the Access Node.

For Ethernet access aggregation, where a per-subscriber (stacked) VLAN can be applied (1:1 model as defined in [TR-101]), the TLV contains two value fields. Each field carries a 12-bit VLAN identifier (which is part of the VLAN tag defined by IEEE 802.1Q). The first field MUST carry the inner VLAN identifier, while the second field MUST carry the outer VLAN identifier.

When the N:1 VLAN model is used, only one VLAN tag is available. For the N:1 model, the Access-Aggregation-Circuit-ID-Binary TLV contains a single value field, which MUST carry the 12-bit VLAN identifier derived from the single available VLAN tag.

In the case of an ATM aggregation network, where the DSLAM is directly connected to the NAS (without an intermediate ATM switch), the VPI and VCI on the DSLAM uplink correspond uniquely to the DSL line on the DSLAM. The Access-Aggregation-Circuit-ID-Binary TLV MAY be used to carry the VPI and VCI. The first value field of the TLV MUST carry the VCI, while the second value field MUST carry the VPI.

Each identifier MUST be placed in the low-order bits of its respective 32-bit field, with the higher-order bits set to zero. The ordering of the bits of the identifer MUST be the same as when the identifier is transmitted on the wire to identify an Ethernet frame or ATM cell.

The Access-Aggregation-Circuit-ID-Binary is illustrated in Figure 13.

Length: 4 or 8 bytes

Value: one or two 32-bit binary fields.

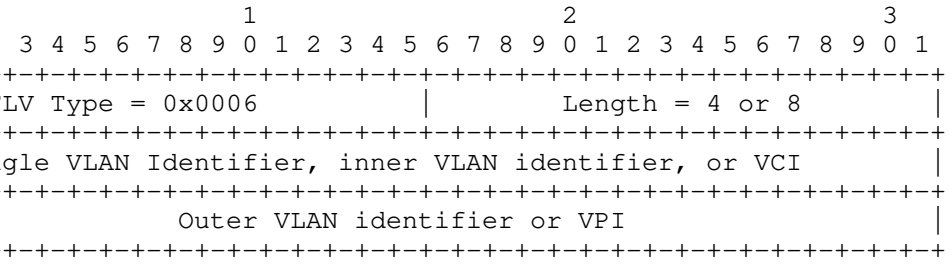

Figure 13: The Access-Aggregation-Circuit-ID-Binary TLV

5.1.2.4. Access-Aggregation-Circuit-ID-ASCII TLV

Type: 0x0003

Description: This TLV transmits the ASCII equivalent of the Access-Aggregation-Circuit-ID-Binary TLV. As mentioned in the previous section, the AN control application will use a format similar to that specified in Section 3.9.3 of [TR-101] for the format of the "circuit-id".

As an extension to the present document, the Access Node could convey to the NAS the characteristics (e.g., bandwidth) of the uplink on the Access Node. This TLV or the binary equivalent defined above then serves the purpose of uniquely identifying the

uplink whose characteristics are being defined. The present document does not specify the TLVs needed to convey the uplink characteristics.

Length: up to 63 bytes

Value: ASCII string

6. ANCP Based DSL Topology Discovery

Section 3.1 of [RFC5851] describes the requirements for the DSL Topology Discovery capability.

6.1. Control Context (Informative)

The AN control application in the DSLAM requests ANCP to send a DSL-specific Port Up message to the NAS under the following circumstances:

- when a new adjacency with the NAS is established, for each DSL loop that is synchronized at that time;
- subsequent to that, whenever a DSL loop resynchronizes; and
- whenever the AN control application wishes to signal that a line attribute has changed.

The AN control application in the DSLAM requests ANCP to send a DSL-specific Port Down message to the NAS under the following circumstances:

- when a new adjacency with the NAS is established, for each DSL loop that is provisioned but not synchronized at that time;
- whenever a DSL loop that is equipped in an AN but administratively disabled is signalled as "IDLE"; and
- subsequent to that, whenever a DSL loop loses synchronization.

The AN control application passes information to identify the DSL loop to ANCP to include in the Port Up or Port Down message, along with information relating to DSL loop attributes.

In the case of bonded copper loops to the customer premise (as per DSL multi-pair bonding described by [G.988.1] and [G.988.2]), the AN control application requests that ANCP send DSL-specific Port Up and Port Down messages for the aggregate "DSL bonded circuit"

(represented as a single logical port) as well as the individual DSL loops of which it is comprised. The information relating to DSL line attributes that is passed by the AN control application is aggregate information.

ANCP generates the DSL-specific Port Up or Port Down message and transfers it to the NAS. ANCP on the NAS side passes an indication to the NAS control application that a DSL Port Up or Port Down message has been received along with the information contained in the message.

The NAS control application updates its view of the DSL loop state, performs any required accounting operations, and uses any included line attributes to adjust the operation of its queueing/scheduling mechanisms as they apply to data passing to and from that DSL loop.

Figure 14 summarizes the interaction.

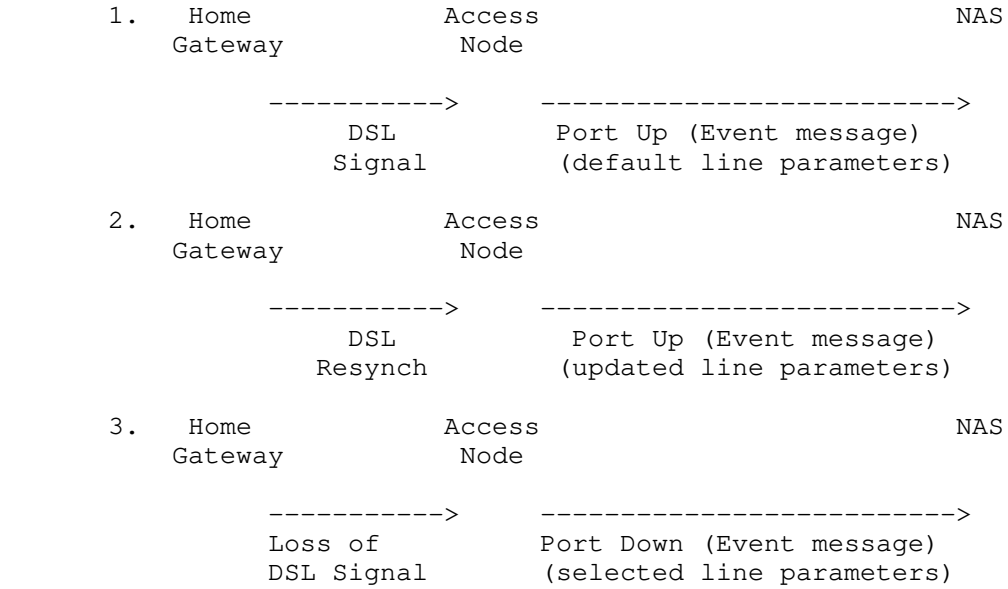

Figure 14: ANCP Message Flow For DSL Topology Discovery

6.2. Protocol Requirements

The DSL topology discovery capability is assigned capability type 0x0001. No capability data is associated with this capability.

6.2.1. Protocol Requirements On the AN Side

The AN-side ANCP agent MUST be able to create DSL-specific Port Up and Port Down messages according to the format specified in Section 6.3.

The AN-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The AN-side ANCP agent MUST follow the AN-side procedures associated with DSL-specific Port Up and Port Down messages as they are specified in Section 6.4.

6.2.2. Protocol Requirements On the NAS Side

The NAS-side ANCP agent MUST be able to receive and validate DSL-specific Port Up and Port Down messages according to the format specified in Section 6.3.

The NAS-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The NAS-side ANCP agent MUST follow the NAS-side procedures associated with DSL-specific Port Up and Port Down messages as they are specified in Section 6.4.

6.3. ANCP Port UP and Port DOWN Event Message Descriptions

The format of the ANCP Port UP and Port DOWN Event messages is shown in Figure 15.

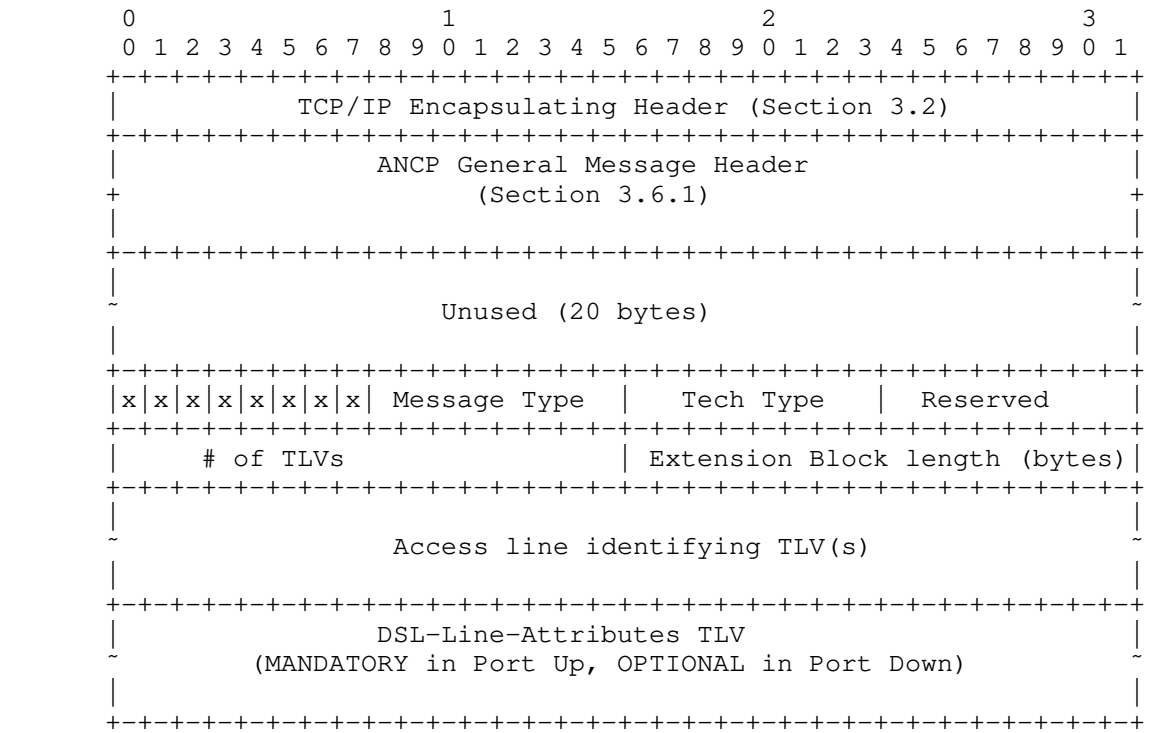

NOTE: TLVs MAY be in a different order from what is shown in this figure.

Figure 15: Format Of the ANCP Port Up and Port Down Event Messages For DSL Topology Discovery

See Section 3.6.1 for a description of the ANCP general message header. The Message Type field MUST be set to 80 for Port Up, 81 for Port Down. The 12 bit Result Code field MUST be set to 0. The 4 bit Result field MUST be set to 0 (signifying Ignore). The 24-bit Transaction Identifier field MUST be set to 0. Other fields in the general header MUST be set as described in Section 3.6.

The five word Unused field is a historical leftover. The handling of unused/reserved fields is described in Section 3.4.

The remaining message fields belong to the "extension block", and are described as follows:

Extension Flags: The flag bits denoted by 'x' are currently unspecified and reserved.

Message Type: Message Type has the same value as in the general header (i.e., 80 or 81).

Tech Type: MUST be set to 0x05 (DSL).

# of TLVs: the number of TLVs that follow, not counting TLVs encapsulated within other TLVs.

Extension Block Length: the total length of the TLVs carried in the extension block in bytes, including any padding within individual TLVs.

TLVs: one or more TLVs to identify a DSL line and zero or more TLVs to define its characteristics.

6.4. Procedures

6.4.1. Procedures On the AN Side

The AN-side ANCP agent creates and transmits a DSL-specific Port Up or Port Down message when requested by the AN control application and presented with the information needed to build a valid message. It is RECOMMENDED that the Access Node use a dampening mechanism per DSL loop to control the rate at which state changes are communicated to the NAS.

At the top level, the extension block within a DSL-specific Port Up or Port Down message MUST include TLVs from Section 5.1.2 to identify the DSL loop.

TLVs presenting DSL line attributes (i.e., the TLVs specified in Section 6.5) MUST be encapsulated within the DSL-Line-Attributes TLV. When the DSL-Line-Attributes TLV is present in a message, it MUST contain at least one such TLV and will generally contain more than one. In the Port Up message, the DSL-Line-Attributes TLV MUST be present. In the Port Down message, the DSL-Line-Attributes TLV MAY be present.

6.4.2. Procedures On the NAS Side

The NAS-side ANCP agent MUST be prepared to receive Port Up and Port Down messages for a given DSL loop or logical port at any time after negotiation of an adjacency has been completed. It is possible for two Port Up messages in succession to be received for the same DSL loop without an intervening Port Down message, and vice versa.

The NAS-side ANCP agent SHOULD validate each message against the specifications given in Section 6.3 and the TLV specifications given in Section 5.1.2 and Section 6.5. If it finds an error it MAY generate a Generic Response message containing an appropriate Result Code value. If it does so, the message MUST contain copies of all of the identifier TLVs from Section 5.1.2 that were present in the Port Up or Port Down message. The message SHOULD also contain a Status-Info TLV which in turn contains other information appropriate to the message header Result Code value as described in Section 3.6.1.4.

If the received message passes validation, the NAS-side ANCP agent extracts the information from the TLVs contained in the message and presents that information along with an indication of reported event type to the NAS control application. If validation of individual TLVs fails but the message as a whole can be processed, the NAS-side ANCP agent "may" pass the valid message contents to the NAS control application.

6.5. TLVs For DSL Line Attributes

As specified above, the DSL-Line-Attributes TLV is inserted into the Port Up or Port Down message at the top level. The remaining TLVs defined below are encapsulated within the DSL-Line-Attributes TLV.

6.5.1. DSL-Line-Attributes TLV

Type: 0x0004

Description: This TLV encapsulates attribute values for a DSL line serving a subscriber.

Length: variable (up to 1024 bytes)

Value: one or more encapsulated TLVs corresponding to DSL line attributes. The DSL-Line-Attributes TLV MUST contain at least one TLV when it is present in a Port Up or Port Down message. The actual contents are determined by the AN control application.

6.5.2. DSL-Type TLV

Type: 0x0091

Description: Indicates the type of transmission system in use.

Length: 4 bytes

Value: 32 bit unsigned integer

ADSL1 = 1

ADSL2 = 2

ADSL2+ = 3

VDSL1 = 4

VDSL2 = 5

SDSL = 6

OTHER = 0

6.5.3. Actual-Net-Data-Rate-Upstream TLV

Type: 0x0081

Description: Actual upstream net data rate on a DSL line.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.4. Actual-Net-Data-Rate-Downstream TLV

Type: 0x0082

Description: Actual downstream net data rate on a DSL line.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.5. Minimum-Net-Data-Rate-Upstream TLV

Type: 0x0083

Description: Minimum upstream net data rate desired by the operator.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.6. Minimum-Net-Data-Rate-Downstream TLV

Type: 0x0084

Description: Minimum downstream net data rate desired by the operator.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.7. Attainable-Net-Data-Rate-Upstream TLV

Type: 0x0085

Description: Maximum net upstream rate that can be attained on the DSL line.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.8. Attainable-Net-Data-Rate-Downstream TLV

Type: 0x0086

Description: Maximum net downstream rate that can be attained on the DSL line.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.9. Maximum-Net-Data-Rate-Upstream TLV

Type: 0x0087

Description: Maximum net upstream data rate desired by the operator.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.10. Maximum-Net-Data-Rate-Downstream TLV

Type: 0x0088

Description: Maximum net downstream data rate desired by the operator.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.11. Minimum-Net-Low-Power-Data-Rate-Upstream TLV

Type: 0x0089

Description: Minimum net upstream data rate desired by the operator in low power state.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.12. Minimum-Net-Low-Power-Data-Rate-Downstream TLV

Type: 0x008A

Description: Minimum net downstream data rate desired by the operator in low power state.

Length: 4 bytes

Value: Rate in Kbits/s as a 32 bit unsigned integer

6.5.13. Maximum-Interleaving-Delay-Upstream TLV

Type: 0x008B

Description: maximum one way interleaving delay.

Length: 4 bytes

Value: Time in ms as a 32 bit unsigned integer

6.5.14. Actual-Interleaving-Delay-Upstream TLV

Type: 0x008C

Description: Value corresponding to the interleaver setting.

Length: 4 bytes

Value: Time in ms as a 32 bit unsigned integer

6.5.15. Maximum-Interleaving-Delay-Downstream TLV

Type: 0x008D

Description: maximum one way interleaving delay.

Length: 4 bytes

Value: Time in ms as a 32 bit unsigned integer

6.5.16. Actual-Interleaving-Delay-Downstream

Type: 0x008E

Description: Value corresponding to the interleaver setting.

Length: 4 bytes

Value: Time in ms as a 32 bit unsigned integer

6.5.17. DSL-Line-State TLV

Type: 0x008F

Description: The state of the DSL line.

Length: 4 bytes

Value: 32 bit unsigned integer

SHOWTIME = 1

IDLE = 2

SILENT = 3

6.5.18. Access-Loop-Encapsulation TLV

Type: 0x0090

Description: The data link protocol and, optionally, the encapsulation overhead on the access loop. When this TLV is present, at least the data link protocol MUST be indicated. The encapsulation overhead MAY be indicated. The Access Node MAY choose to not convey the encapsulation on the access loop by specifying values of 0 (NA) for the two encapsulation fields.

Length: 3 bytes

Value: The three bytes (most to least significant) and valid set of values for each byte are defined as follows:

Byte 1: Data Link

ATM AAL5 = 0

ETHERNET = 1

Byte 2: Encapsulation 1

NA = 0

Untagged Ethernet = 1

Single-tagged Ethernet = 2

Double-tagged Ethernet = 3

Byte 3: Encapsulation 2

NA = 0

PPPoA LLC = 1

PPPoA NULL = 2

IPoA LLC = 3

IPoA NuLL = 4

Ethernet over AAL5 LLC with FCS = 5

Ethernet over AAL5 LLC without FCS = 6

Ethernet over AAL5 NULL with FCS = 7

Ethernet over AAL5 NULL without FCS = 8

The Access-Loop-Encapsulation TLV is illustrated in Figure 16.

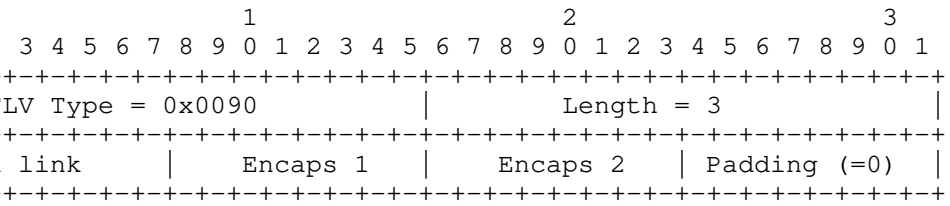

Figure 16: The Access-Loop-Encapsulation TLV

7. ANCP based DSL Line Configuration

The use case for ANCP-based DSL Line Configuration is described in Section 3.2 of [RFC5851].

7.1. Control Context (Informative)

Triggered by topology information reporting a new DSL line or triggered by a subsequent user session establishment (via PPP or DHCP), RADIUS/AAA sends service parameters to the NAS control application for configuration on the access line. The NAS control application passes the request on to the NAS-side agent, which sends the information to the AN by means of a Port Management (line configuration) message. The AN-side agent passes this information up to the AN control application, which applies it to the line. Figure 17 summarizes the interaction.

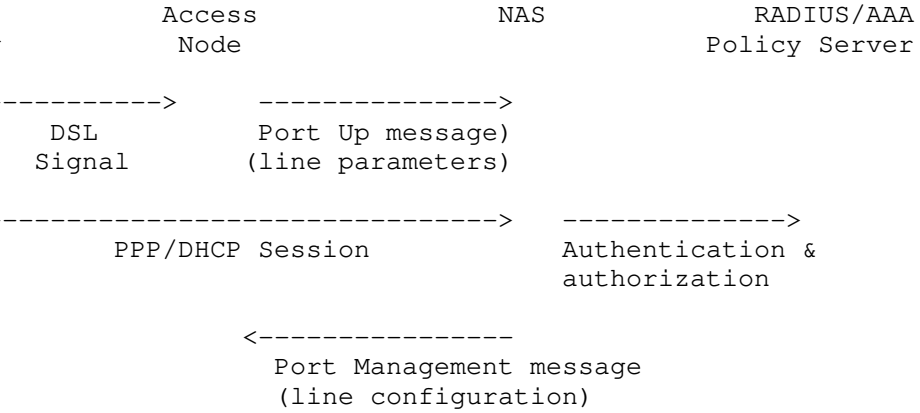

Figure 17: Message Flow - ANCP Mapping For Initial Line Configuration

The NAS could update the line configuration as a result of a subscriber service change (e.g. triggered by the policy server). Figure 18 summarizes the interaction.

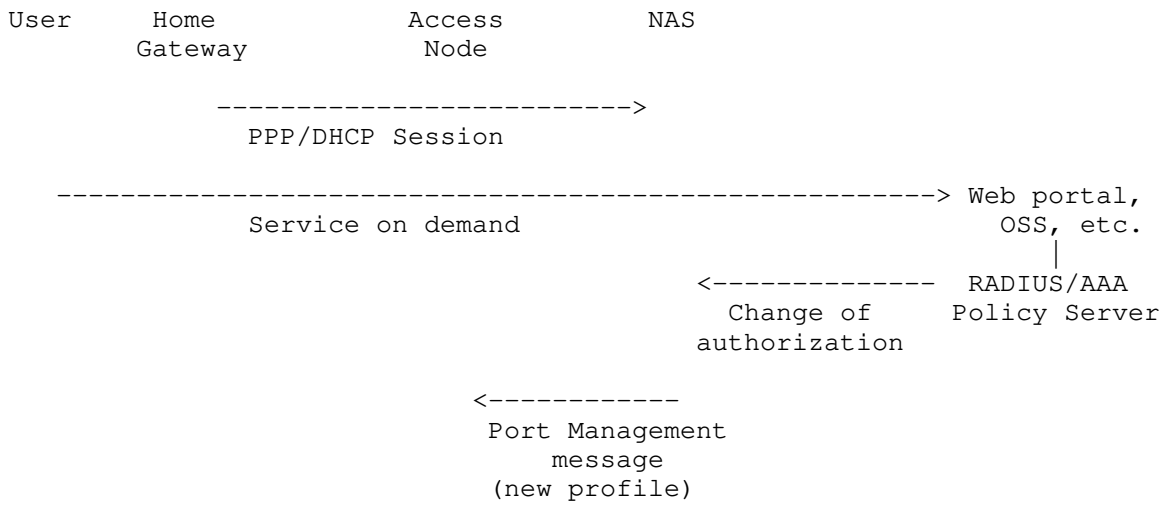

Figure 18: Message flow - ANCP Mapping For Updated Line Configuration

7.2. Protocol Requirements

The DSL line configuration capability is assigned capability type 0x0002. No capability data is associated with this capability.

7.2.1. Protocol Requirements On the NAS Side

The NAS-side ANCP agent MUST be able to create DSL-specific Port Management (line configuration) messages according to the format specified in Section 7.3.

The NAS-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The NAS-side ANCP agent MUST follow the NAS-side procedures associated with DSL-specific Port Management (line configuration) messages as they are specified in Section 7.4.

7.2.2. Protocol Requirements On the AN Side

The AN-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The AN-side ANCP agent MUST be able to receive and validate DSL-

specific Port Management (line configuration) messages according to the format specified in Section 7.3.

The AN-side ANCP agent MUST follow the AN-side procedures associated with DSL-specific Port Management (line configuration) messages as specified in Section 7.4.

7.3. ANCP Port Management (Line Configuration) Message Format

The ANCP Port Management message for DSL line configuration has the format shown in Figure 19.

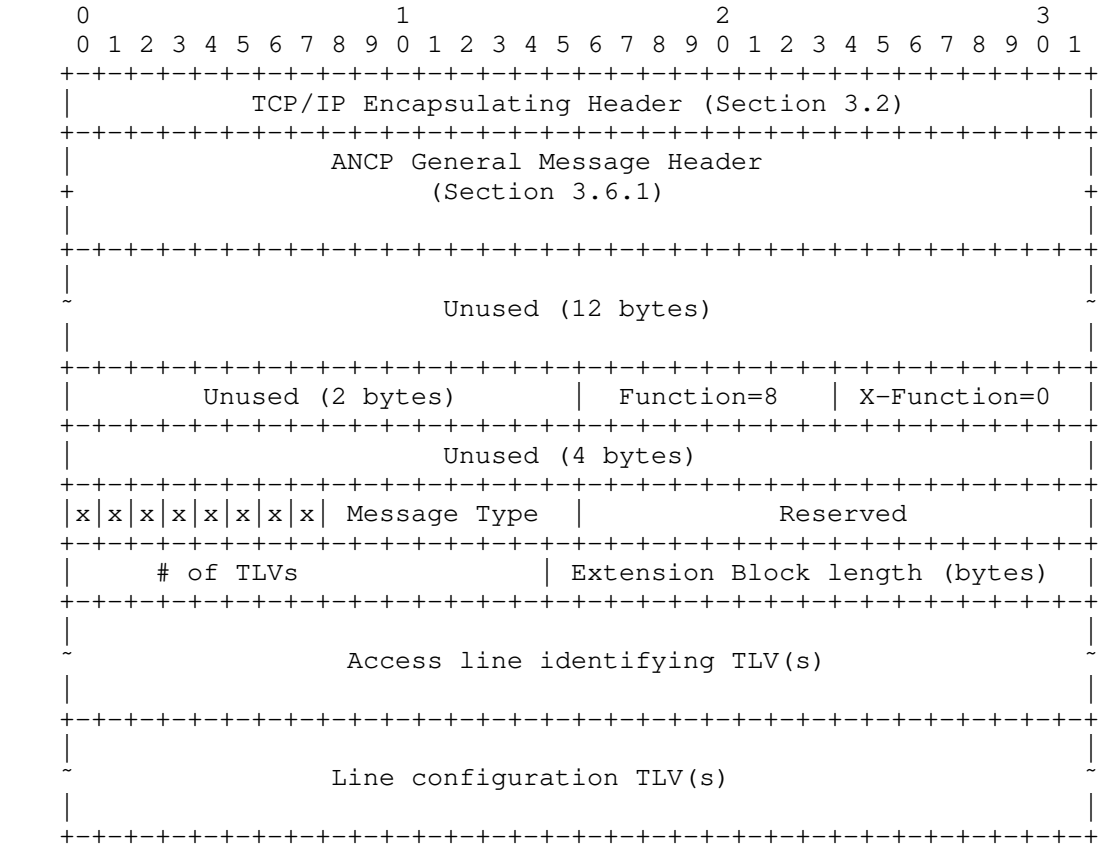

NOTE: TLVs MAY be in a different order from what is shown in this figure.

Figure 19: Port Management Message For DSL Line Configuration

See Section 3.6 for a description of the ANCP general message header.

The Message Type field MUST be set to 32. The 12 bit Result Code field MUST be set to 0x0. The 4 bit Result field MUST be set to either 1 (NAck) or 2 (AckAll), as determined by policy on the NAS. The 24-bit Transaction Identifier field MUST be set to a positive value. Other fields in the general header MUST be set as described in Section 3.6.

The handling of the various unused/reserved fields is described in Section 3.4.

The remaining message fields are described as follows:

Function: action to be performed. For line configuration, Function MUST be set to 8 (Configure Connection Service Data). This action type requests the Access Node (i.e., DSLAM) to apply service configuration data contained in the line configuration TLVs to the DSL line designated by the access line identifying TLVs.

X-Function: qualifies the action set by Function. For DSL line configuration, this field MUST be set to 0.

Extension Flags: the flag bits denoted by 'x' before the Message Type field are reserved for future use.

Message Type: Message Type has the same value as in the general header (i.e., 32).

Reserved (16 bits): reserved for future use.

# of TLVs: the number of TLVs that follow, not counting TLVs encapsulated within other TLVs.

Extension Block Length: the total length of the TLVs carried in the extension block in bytes, including any padding within individual TLVs.

TLVs: two or more TLVs to identify a DSL line and configure its service data.

Other ANCP capabilities, either specific to DSL or technology-independent, MAY reuse the Port Management message for service configuration. If the settings of the fixed fields are compatible with the settings just described, the same Port Management message that is used for DSL line configuration MAY be used to carry TLVs relating to the other capabilities that apply to the same DSL loop.

Use of the Port Management message for configuration MAY also be generalized to other access technologies, if the respective

capabilities specify use of access line identifiers appropriate to those technologies in place of the identifiers defined in Section 5.1.2.

7.4. Procedures

Service configuration MAY be performed on an access line regardless of its current state.

7.4.1. Procedures On the NAS Side

When requested by the NAS control application and presented with the necessary information to do so, the NAS-side agent MUST create and send a Port Management message with the fixed fields set as described in the previous section. The message MUST contain one or more TLVs to identify an access line according the requirements of Section 5.1.2. The NAS MUST include one or more TLVs to configure line service parameters for that line. Section 7.5 currently identifies only one such TLV, Service-Profile-Name, but other TLVs MAY be added by extensions to ANCP.

7.4.2. Procedures On the AN Side

The AN-side ANCP agent MUST be prepared to receive Port Management (line configuration) messages for a given DSL loop or logical port at any time after negotiation of an adjacency has been completed.

The AN-side ANCP agent SHOULD validate each message against the specifications given in Section 7.3 and the TLV specifications given in Section 5.1.2 and Section 7.5. If it finds an error it MUST return a Port Management response message which copies the Port Management request as it was received, but has the Result header field set to 0x04 (Failure) and the Result Code field set to the appropriate value. The AN-side agent MAY add a Status-Info TLV (Section 4.5) to provide further information on the error, particularly if this is recommended in Section 3.6.1.4 for the given Result Code value. If it does so, the various length fields and the # of TLVs field within the message MUST be adjusted accordingly.

If the received message passes validation, the AN-side ANCP agent "must" extract the information from the TLVs contained in the message and present that information to the AN control application. In addition, if the Result header field was set to 0x2 (AckAll) in the original request, the AN-side agent "must" indicate to the AN control application that a response is required. When the AN control application indicates that it has processed the request successfully, the AN-side agent MUST return a Port Management response message which duplicates the request except that the Result header field is

set to 0x3 (Success). (The Result Code field, as in the original request, has value 0.)

7.5. TLVs For DSL Line Configuration

Currently only the following TLV is specified for DSL line configuration. More TLVs may be defined in a future version of this specification or in ANCP extensions for individual service attributes of a DSL line (e.g. rates, interleaving delay, multicast channel entitlement access-list).

7.5.1. Service-Profile-Name TLV

Type: 0x0005

Description: Reference to a pre-configured profile on the DSLAM that contains service specific data for the subscriber.

Length: up to 64 bytes

Value: ASCII string containing the profile name (which the NAS learns from a policy server after a subscriber is authorized).

8. ANCP-Based DSL Remote Line Connectivity Testing

The use case and requirements for ANCP-Based DSL remote line connectivity testing are specified in Section 3.3 of [RFC5851]

8.1. Control Context (Informative)

The NAS control application initiates a request for remote connectivity testing for a given access loop. The NAS control application can provide loop count and timeout test parameters and opaque data for its own use with the request. The loop count parameter indicates the number of test messages or cells to be used. The timeout parameter indicates the longest that the NAS control application will wait for a result.

The request is passed in a Port Management (OAM) message. If the NAS control application has supplied test parameters, they are used, otherwise the AN control application uses default test parameters. If a loop count parameter provided by the NAS is outside the valid range, the AN does not execute the test, but returns a result indicating that the test has failed due to an invalid parameter. If the test takes longer than the timeout value (default or provided by the NAS) the AN control application can return a failure result indicating timeout or else can send no response. The AN control

application can provide a human-readable string describing the test results, for both failures and successes. If provided, this string is included in the response. Responses always include the opaque data, if any, provided by the NAS control application.

Figure 20 summarizes the interaction.

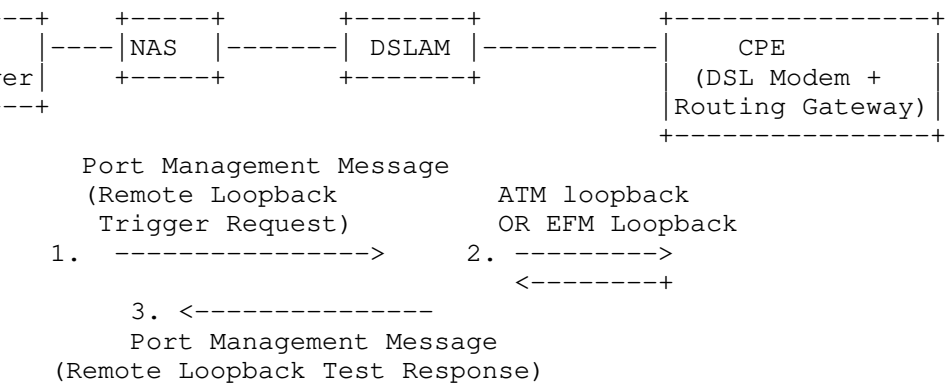

Figure 20: Message Flow For ANCP based OAM

8.2. Protocol Requirements

The DSL remote line connectivity testing capability is assigned capability type 0x0004. No capability data is associated with this capability.

8.2.1. Protocol Requirements On the NAS Side

The NAS-side ANCP agent MUST be able to create DSL-specific Port Management (OAM) messages according to the format specified in Section 8.3.

The NAS-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The NAS-side ANCP agent MUST follow the NAS-side procedures associated with DSL-specific Port Management (OAM) messages as they are specified in Section 8.4.

8.2.2. Protocol Requirements On the AN Side

The AN-side ANCP agent MUST conform to the normative requirements of Section 5.1.2.

The AN-side ANCP agent MUST be able to receive and validate DSL-

specific Port Management (OAM) messages according to the format specified in Section 8.3.

The AN-side ANCP agent MUST follow the AN-side procedures associated with DSL-specific Port Management (OAM) messages as specified in Section 8.4.

8.3. Port Management (OAM) Message Format

The Port Management message for DSL line testing has the same format as for DSL line configuration (see Section 7.3), with the following differences:

- The Result field in the request SHOULD be set to AckAll (0x1), to allow the NAS to receive the information contained in a successful test response.
- The Function field MUST be set to 9 (Remote Loopback). (The X-Function field continues to be 0.)
- The appended TLVs in the extension value field include testing-related TLVs rather than subcriber service information.

The Port Management (OAM) message is illustrated in Figure 21.

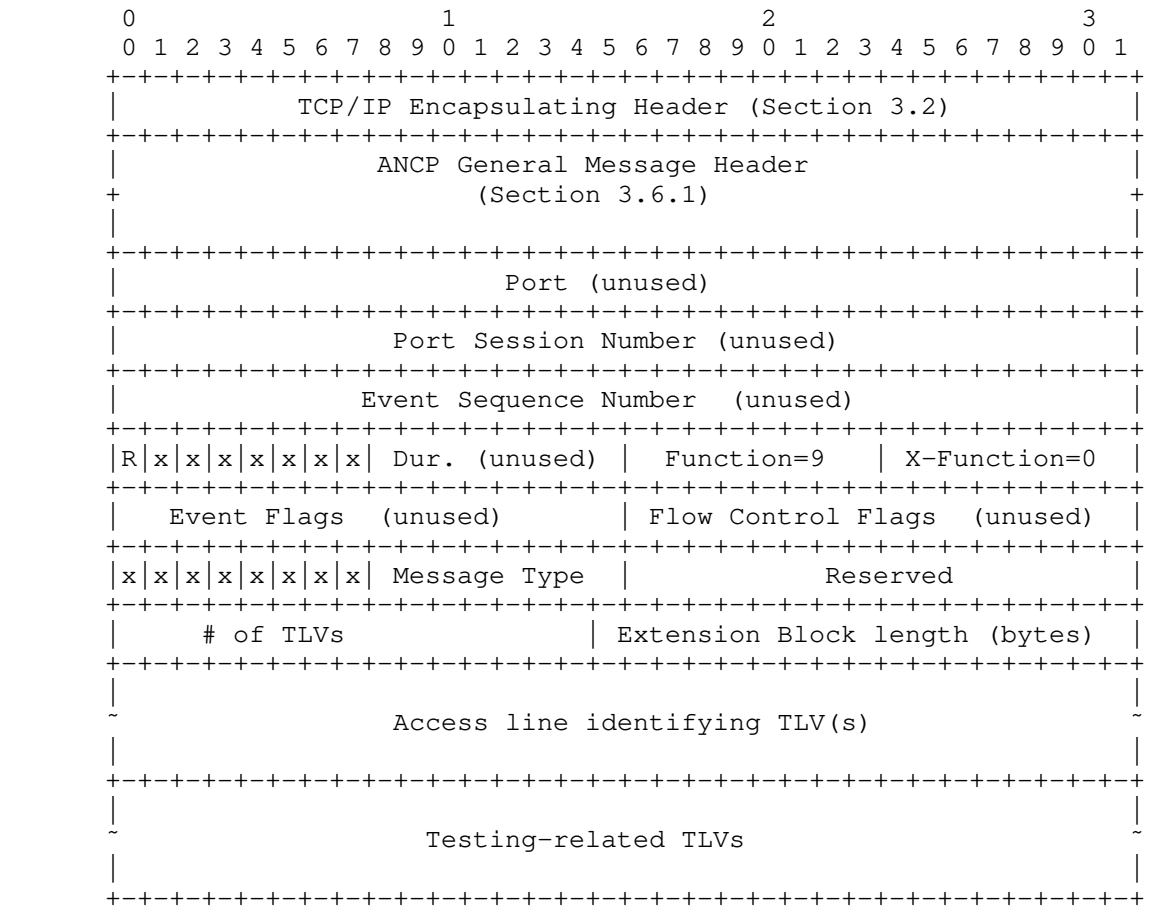

NOTE: TLVs MAY be in a different order from what is shown in this figure.

Figure 21: Port Management Message For DSL Line Remote Connectivity Testing

8.4. Procedures

From the point of view of ANCP, it is permissible to attempt line connectivity testing regardless of the state of the line. However, testing could fail in some states due to technology limitations.

8.4.1. NAS-Side Procedures

When requested by the NAS control application and presented with the necessary information to do so, the NAS-side agent creates and sends

a Port Management (OAM) request with the fixed fields set as described in the previous section. The message MUST contain one or more TLVs to identify an access line according the requirements of Section 5.1.2. The NAS MAY include the Opaque-Data TLV and/or the OAM-Loopback-Test-Parameters TLV (defined in Section 8.5) to configure the loopback test for that line.

8.4.2. AN-Side Procedures

The AN-side ANCP agent SHOULD validate each message against the specifications given in Section 8.3 and the TLV specifications given in Section 5.1.2 and Section 8.5. If it finds an error it MUST return a Port Management response message which copies the Port Management request as it was received, but has the Result header field set to 0x04 (Failure) and the Result Code field set to the appropriate value. Result Code value 0x509 as described below MAY apply, as well as the other Result Code values documented in Section 3.6.1.4. Result Code value 0x509 SHOULD be used if the OAM-Loopback-Test-Parameters TLV is present with an invalid value of the Count field. The AN-side agent MAY add a Status-Info TLV (Section 4.5) to provide further information on the error, particularly if this is recommended in Section 3.6.1.4 for the given Result Code value. If it does so, the various length fields and the # of TLVs field within the message MUST be adjusted accordingly.

If the received message passes validation, the AN-side ANCP agent extracts the information from the TLVs contained in the message and presents that information to the AN control application. It MUST NOT generate an immediate response to the request, but MUST instead wait for the AN control application to indicate that the response should be sent.

When requested by the AN control application and presented with the necessary information to do so, the AN-side agent creates and sends a Port Management (OAM) response to the original request. The Result field MUST be set to Success (0x3) or Failure (0x4), and the Result Code field SHOULD be set to one of the following values, as indicated by the AN control application.

0x500: Specified access line does not exist. See the documentation of Result Code 0x500 in Section 3.6.1.4 for more information. The Result header field MUST be set to Failure (0x4).

0x501: Loopback test timed out. The Result header field MUST be set to Failure (0x4).

0x503: DSL line status showtime

0x504: DSL line status idle

0x505: DSL line status silent

0x506: DSL line status training

0x507: DSL line integrity error

0x508: DSLAM resource not available. The Result header field MUST be set to Failure (0x04).

0x509: Invalid test parameter. The Result header field MUST be set to Failure (0x4).

All other fields of the request including the TLVs MUST be copied into the response unchanged, except that in a successful response the OAM-Loopback-Test-Parameters TLV MUST NOT appear. If the AN control application has provided the necessary information, the AN-side agent MUST also include an instance of the OAM-Loopback-Test-Response-String TLV in the response.

8.5. TLVs For the DSL Line Remote Connectivity Testing Capability

The following TLVs have been defined for use with the DSL line testing capability.

8.5.1. OAM-Loopback-Test-Parameters TLV

Type: 0x0007

Description: Parameters intended to override the default values for this loopback test.

Length: 2 bytes

Value: two unsigned 1 byte fields described below (listed in order of most to least significant).

Byte 1: Count. Number of loopback cells/messages that should be generated on the local loop as part of the loopback test. The Count value SHOULD be greater than 0 and less than or equal to 32.

Byte 2: Timeout. Upper bound on the time in seconds that the NAS will wait for a response from the DSLAM. The value 0 MAY be used to indicate that the DSLAM MUST use a locally

determined value for the timeout.

The OAM-Loopback-Test-Parameters TLV is illustrated in Figure 22

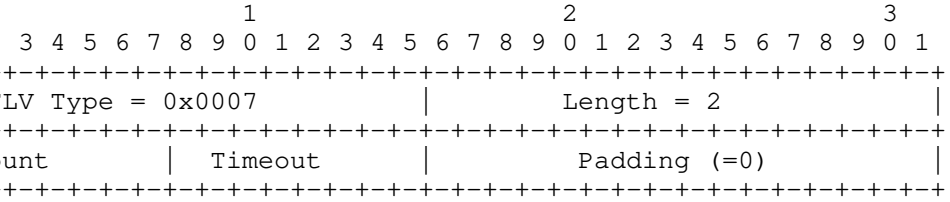

Figure 22: The OAM-Loopback-Test-Parameters TLV

8.5.2. Opaque-Data TLV

Type: 0x0008

Description: An 8 byte opaque field used by the NAS control application for its own purposes (e.g., response correlation.) The procedures in Section 8.4.2 ensure that if it is present in the request it is copied unchanged to the response.

Length: 8 bytes

Value: Two 32 bit unsigned integers.

8.5.3. OAM-Loopback-Test-Response-String TLV

Type: 0x0009

Description: Suitably formatted string containing useful details about the test that the NAS will display for the operator, exactly as received from the DSLAM (no manipulation or interpretation by the NAS).

Length: up to 128 bytes

Value: UTF-8 encoded string of text.

9. IANA Considerations

IANA NOTE: please replace "RFCXXXX" with the number of this specification.

9.1. Summary

This section requests the following IANA actions:

- establishment of the following new ANCP registries:

  ANCP Message Types;

  ANCP Result Codes;

  ANCP Port Management Functions;

  ANCP Technology Types;

  ANCP Command Codes;

  ANCP TLV Types;

  ANCP Capabilities.

- establishment of a new joint GSMP/ANCP version registry;

- addition of ANCP as another user of TCP port 6068 in the port number registry at http://www.iana.org/assignments/port-numbers. The current user is GSMP.

All of these actions are described in detail below except for the port registration, for which the final point above should provide sufficient information.

9.2. IANA Actions

9.2.1. ANCP Message Type Registry

IANA is requested to create a new registry, Access Network Control Protocol (ANCP) Message Types. Additions to that registry are permitted by Standards Action, as defined by [RFC5226]. The values for Message Type MAY range from 0 to 255, but new Message Types SHOULD be assigned values sequentially from 90 onwards (noting that 91 and 93 are already assigned). The initial contents of the ANCP Message Types registry are as follows:

Message Type	Message Name	Reference
10	Adjacency Protocol	RFCXXXX
32	Port Management	RFCXXXX
80	Port Up	RFCXXXX
81	Port Down	RFCXXXX
85	Adjacency Update	RFCXXXX
91	Generic Response	RFCXXXX
93	Provisioning	RFCXXXX

9.2.2. ANCP Result Code Registry

IANA is requested to create a new registry, Access Network Control Protocol (ANCP) Result Codes. The documentation of new Result Codes MUST include the following information:

- Result Code value TBD (as assigned by IANA);
- One-line description;
- Where condition detected: (control application or ANCP agent);
- Further description (if any);
- Required additional information in the response message;
- Target (control application or ANCP agent at the peer that sent the original request);
- Action RECOMMENDED for the receiving ANCP agent

The values for Result Code are expressed in hexadecimal, and MAY range from 0x0 to 0xFFFFFF. The range 0x0 to 0xFFF is reserved for allocation by the criterion of IETF Review, as defined by [RFC5226]. IANA SHOULD allocate new Result Code values from this range sequentially beginning at 0x100. The range 0x1000 onwards is allocated by the criterion of Specification Required, as defined by [RFC5226]. IANA SHOULD allocate new Result Code values from this range sequentially beginning at 0x1000. The initial contents of the ANCP Message Types registry are as follows:

Result Code	One-line description	Reference
0x0	No result	RFCXXXX
0x2	Invalid request message	RFCXXXX
0x6	One or more of the specified ports are down	RFCXXXX
0x13	Out of resources	RFCXXXX
0x51	Request message type not implemented	RFCXXXX
0x53	Malformed message	RFCXXXX
0x54	Mandatory TLV missing	RFCXXXX
0x55	Invalid TLV contents	RFCXXXX
0x500	One or more of the specified ports do not exist	RFCXXXX
0x501	Loopback test timed out (0x501)	RFCXXXX
0x502	Reserved (0x502)	RFCXXXX
0x503	DSL line status showtime (0x503)	RFCXXXX
0x504	DSL line status idle (0x504)	RFCXXXX
0x505	DSL line status silent (0x505)	RFCXXXX
0x506	DSL line status training (0x506)	RFCXXXX
0x507	DSL line integrity error (0x507)	RFCXXXX
0x508	DSLAM resource not available (0x508)	RFCXXXX
0x509	Invalid test parameter (0x509)	RFCXXXX

9.2.3. ANCP Port Management Function Registry

IANA is requested to create a new Access Network Control Protocol (ANCP) Port Management Function registry, with the following initial entries. Additions to this registry will be by Standards Action, as defined by [RFC5226]. Values may range from 0 to 255. IANA SHOULD assign values sequentially beginning with 1, taking account of the values already assigned below.

> NOTE: future extensions of ANCP may need to establish sub-registries of permitted X-Function values for specific values of Function.

Function Value	Function Name	Reference
0	Reserved	RFCXXXX
8	Configure Connection Service Data	RFCXXXX
9	Remote Loopback	RFCXXXX

9.2.4. ANCP Technology Type Registry

IANA is requested to create a new Access Network Control Protocol (ANCP) Technology Type registry, with additions by Expert Review, as defined by [RFC5226]. The Technology Type MUST designate a distinct access transport technology. Values may range from 0 to 255. IANA SHOULD assign new values sequentially beginning at 2, taking into account of the values already assigned below. The initial entries are as follows:

Tech Type Value	Tech Type Name	Reference
0	Not technology dependent	RFCXXXX
1	PON	RFCXXXX
5	DSL	RFCXXXX
255	Reserved	RFCXXXX

9.2.5. ANCP Command Code Registry

IANA is requested to create a new Access Network Control Protocol (ANCP) Command Code registry, with additions by Standards Action, as defined by [RFC5226]. Values may range from 0 to 255. IANA SHOULD assign new values sequentially beginning with 1. The initial entry is as follows:

Command Code Value	Command Code Directive Name	Reference
0	Reserved	RFCXXXX

9.2.6. ANCP TLV Type Registry

IANA is requested to create a new Access Network Control Protocol (ANCP) TLV Type registry. Values are expressed in hexadecimal and may range from 0x0000 to 0xFFFF. Additions in the range 0x0000 to 0x1FFF are by IETF Review, as defined by [RFC5226]. IANA SHOULD assign new values in this range sequentially beginning at 0x100 and taking account of the assignments already made below. Additions in the range 0x2000 to 0xFFFF are by Specification Required, again as defined by [RFC5226]. IANA SHOULD assign new values in this range sequentially beginning at 0x2000. In both cases, the documentation of the TLV MUST provide:

- a TLV name following the convention used for the initial entries (capitalized words separated by hyphens);

- a brief description of the intended use;
- a precise description of the contents of each fixed field, including its length, type, and units (if applicable);
- identification of any mandatory encapsulated TLVs;
- an indication of whether optional TLVs may be encapsulated, with whatever information is available on their identity (could range from a general class of information to specific TLV names, depending on the nature of the TLV being defined.

The initial entries are as follows:

Type Code	TLV Name	Reference
0x0000	Reserved	RFCXXXX
0x0001	Access-Loop-Circuit-ID	RFCXXXX
0x0002	Access-Loop-Remote-Id	RFCXXXX
0x0003	Access-Aggregation-Circuit-ID-ASCII	RFCXXXX
0x0004	DSL-Line-Attributes	RFCXXXX
0x0005	Service-Profile-Name	RFCXXXX
0x0006	Access-Aggregation-Circuit-ID-Binary	RFCXXXX
0x0007	OAM-Loopback-Test-Parameters	RFCXXXX
0x0008	Opaque-Data	RFCXXXX
0x0009	OAM-Loopback-Test-Response-String	RFCXXXX
0x0011	Command	RFCXXXX
0x0081	Actual-Net-Data-Upstream	RFCXXXX
0x0082	Actual-Net-Data-Rate-Downstream	RFCXXXX
0x0083	Minimum-Net-Data-Rate-Upstream	RFCXXXX
0x0084	Minimum-Net-Data-Rate-Downstream	RFCXXXX
0x0085	Attainable-Net-Data-Rate-Upstream	RFCXXXX
0x0086	Attainable-Net-Data-Rate-Downstream	RFCXXXX
0x0087	Maximum-Net-Data-Rate-Upstream	RFCXXXX
0x0088	Maximum-Net-Data-Rate-Downstream	RFCXXXX
0x0089	Minimum-Net-Low-Power-Data-Rate-Upstream	RFCXXXX
0x008A	Minimum-Net-Low-Power-Data-Rate-Downstream	RFCXXXX
0x008B	Maximum-Interleaving-Delay-Upstream	RFCXXXX
0x008C	Actual-Interleaving-Delay-Upstream	RFCXXXX
0x008D	Maximum-Interleaving-Delay-Downstream	RFCXXXX
0x008E	Actual-Interleaving-Delay-Downstream	RFCXXXX
0x008F	DSL-Line-State	RFCXXXX
0x0090	Access-Loop-Encapsulation	RFCXXXX
0x0091	DSL-Type	RFCXXXX
0x0106	Status-Info	RFCXXXX
0x1000	Target (single access line variant)	RFCXXXX

0x1001 - 0x1020	Reserved for Target variants	RFCXXXX

9.2.7. ANCP Capability Type Registry

IANA is requested to create a new Access Network Control Protocol (ANCP) Capability Type registry, with additions by Standards Action as defined by [RFC5226]. Values may range from 0 to 255. IANA SHOULD assign values sequentially beginning at 5. The specification for a given capability MUST indicate the Technology Type value with which it is associated. The specification MUST further indicate whether the capability is associated with any capability data. Normally a capability is expected to be defined in the same document that specifies the implementation of that capability in protocol terms. The initial entries in the ANCP capability registry are as follows:

Value	Capability Type Name	Tech Type	Capability Data?	Reference
0	Reserved			RFCXXXX
1	DSL Topology Discovery	5	No	RFCXXXX
2	DSL Line Configuration	5	No	RFCXXXX
3	Reserved			RFCXXXX
4	DSL Line Testing	5	No	RFCXXXX

9.2.8. Joint GSMP / ANCP Version Registry

IANA is requested to create a new joint GSMP / ANCP Version registry. Additions to this registry are by Standards Action as defined by [RFC5226]. Values may range from 0 to 255. Values for the General Switch Management Protocol (GSMP) MUST be assigned sequentially beginning with 4 for the next version. Values for the Access Network Control Protocol (ANCP) MUST be assigned sequentially beginning with 50 for the present version. The initial entries are as follows:

Version	Description	Reference
1	GSMP Version 1	RFC1987
2	GSMP Version 2	RFC2297
3	GSMP Version 3	RFC3292
50	ANCP Version 1	RFCXXXX

10. Security Considerations

Security of the ANCP protocol is discussed in [RFC5713]. A number of security requirements on ANCP are stated in Section 8 of that document. Those applicable to ANCP itself are copied to the present document:

- The protocol solution MUST offer authentication of the AN to the NAS.
- The protocol solution MUST offer authentication of the NAS to the AN.
- The protocol solution MUST allow authorization to take place at the NAS and the AN.
- The protocol solution MUST offer replay protection.
- The protocol solution MUST provide data-origin authentication.
- The protocol solution MUST be robust against denial-of-service (DoS) attacks. In this context, the protocol solution MUST consider a specific mechanism for the DoS that the user might create by sending many IGMP messages.
- The protocol solution SHOULD offer confidentiality protection.
- The protocol solution SHOULD ensure that operations in default configuration guarantees a low number of AN/NAS protocol interactions.

Most of these requirements relate to secure transport of ANCP. Robustness against denial-of-service attacks partly depends on transport and partly on protocol design. Ensuring a low number of AN/NAS protocol interactions in default mode is purely a matter of protocol design.

For secure transport, either the combination of IPsec with IKEv2 (references below) or the use of TLS [RFC5246] will meet the requirements listed above. However, the use of TLS has been rejected. The deciding point is a detail of protocol design that was unavailable when [RFC5713] was written. The ANCP adjacency is a major point of vulnerability for denial-of-service attacks. If the adjacency can be shut down, either the AN clears its state pending reestablishment of the adjacency, or the possibility of mismatches between the AN's and NAS's view of state on the AN is opened up. Two ways to cause an adjacency to be taken down are to modify messages so that the ANCP agents conclude that they are no longer synchronized,

or to attack the underlying TCP session. TLS will protect message contents, but not the TCP connection. One has to use either IPsec or the TCP authentication option [RFC5925] for that. Hence the conclusion that ANCP MUST run over IPsec with IKEv2 for authentication and key management.

In greater detail: the ANCP stack MUST include IPsec [RFC4301] running in transport mode, since the AN and NAS are the endpoints of the path. The Encapsulating Security Payload (ESP) [RFC4303] MUST be used, in order to satisfy the requirement for data confidentiality. ESP MUST be configured for the combination of confidentiality, integrity, anti-replay capability. The traffic flow confidentiality service of ESP is unnecessary and, in fact, unworkable in the case of ANCP.

IKEv2 [RFC5996] is also REQUIRED, to meet the requirements for mutual authentication and authorization. Since the NAS and AN MAY be in different trust domains, the use of certificates for mutual authentication could be the most practical approach. However, this is up to the operator(s) concerned.

The AN MUST play the role of initiator of the IKEv2 conversation.

11. Acknowledgements

The authors would like to thank everyone who provided comments or inputs to this document. Swami Subramanian was an early member of the authors' team. The ANCP Working Group is grateful to Roberta Maglione, who served as design team member and primary editor of this document for two years before stepping down. The authors acknowledge the inputs provided by Wojciech Dec, Peter Arberg, Josef Froehler, Derek Harkness, Kim Hyldgaard, Sandy Ng, Robert Peschi, and Michel Platnic, and the further comments provided by Mykyta Yevstifeyev, Brian Carter, Ben Campbell, Alexey Melnikov, Adrian Farrel, Robert Sparks, Peter St. Andre, Sean Turner, and Dan Romascanu.

12. References

12.1. Normative References

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

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

[RFC3293] Worster, T., Doria, A., and J. Buerkle, "General Switch Management Protocol (GSMP) Packet Encapsulations for Asynchronous Transfer Mode (ATM), Ethernet and Transmission Control Protocol (TCP)", RFC 3293, June 2002.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4303] Kent, S., "IP Encapsulating Security Payload (ESP)", RFC 4303, December 2005.

[RFC5646] Phillips, A. and M. Davis, "Tags for Identifying Languages", BCP 47, RFC 5646, September 2009.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

12.2. Informative References

[G.988.1] "ITU-T recommendation G.998.1, ATM-based multi-pair bonding", 2005.

[G.988.2] "ITU-T recommendation G.998.2, Ethernet-based multi-pair bonding,", 2005.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

[RFC3046] Patrick, M., "DHCP Relay Agent Information Option", RFC 3046, January 2001.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[RFC4649] Volz, B., "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option", RFC 4649, August 2006.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security

(TLS) Protocol Version 1.2", RFC 5246, August 2008.

[RFC5713] Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", RFC 5713, January 2010.

[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.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[TR-058] Elias, M. and S. Ooghe, "DSL Forum TR-058, Multi-Service Architecture & Framework Requirements", September 2003.

[TR-059] Anschutz, T., "DSL Forum TR-059, DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services", September 2003.

[TR-092] DSL Forum (now the Broadband Forum), "DSL Forum TR-092, Broadband Remote access server requirements document", 2005.

[TR-101] Cohen et al, "Architecture & Transport: "Migration to Ethernet Based DSL Aggregation", DSL Forum TR-101", 2005.

[TR-147] Voight et al, "Layer 2 Control Mechanism For Broadband Multi-Service Architectures", 2008.

[US_ASCII] American National Standards Institute, "Coded Character Set - 7-bit American Standard Code for Information Interchange", ANSI X.34, 1986.

Authors' Addresses

Sanjay Wadhwa
Alcatel-Lucent

Phone:
Fax:
Email: sanjay.wadhwa@alcatel-lucent.com

Jerome Moisand
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA

Phone:
Fax:
Email: jmoisand@juniper.net

Thomas Haag
Deutsche Telekom
Heinrich-Hertz-Strasse 3-7
Darmstadt, 64295
Germany

Phone: +49 6151 628 2088
Fax:
Email: haagt@telekom.de

Norbert Voigt
Nokia Siemens Networks
Siemensallee 1
Greifswald 17489
Germany

Email: norbert.voigt@nsn.com

Tom Taylor (editor)
Huawei Technologies
Ottawa
Canada

Email: tom111.taylor@bell.net