

Working Group Name  
Internet Draft  
Expires: August 25, 2007

Sanjay Wadhwa  
Juniper Networks

Jerome Moisand  
Juniper Networks

Swami Subramanian  
Juniper Networks

Thomas Haag  
T-Systems

Norbert Voigt  
Siemens

February 25, 2007

## **Protocol for Access Node Control Mechanism in Broadband Networks**

[draft-ietf-ancp-protocol-00.txt](#)

### Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

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

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

This Internet-Draft will expire on August 22, 2007.



## Copyright Notice

Copyright (C) The Internet Society (2007). All Rights Reserved.

## Abstract

This document describes proposed extensions to the GSMPv3 protocol to allow its use in a broadband environment, as a control plane between Access Nodes (e.g. DSLAM) and Broadband Network Gateways (e.g. NAS). These proposed extensions are required to realize a protocol for Access Node Control mechanism as described in [9]. The resulting protocol with the proposed extensions to GSMPv3 [4] is referred to as Access Node Control Protocol (ANCP). This document focuses on specific use cases of access node control mechanism for topology discovery, line configuration, and OAM. It is intended to be augmented by additional protocol specification for future use cases considered in scope by the ANCP charter.

## Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [1].

## Table of Contents

1	Specification Requirements	4
2	Introduction	4
3	Broadband Access Aggregation	4
<a href="#">3.1</a>	ATM-based broadband aggregation.....	<a href="#">4</a>
<a href="#">3.2</a>	Ethernet-based broadband aggregation.....	<a href="#">6</a>
4	Access Node Control Protocol	7
<a href="#">4.1</a>	Overview.....	<a href="#">7</a>
<a href="#">4.2</a>	ANCP based Access Topology Discovery.....	<a href="#">8</a>
<a href="#">4.2.1</a>	Goals.....	<a href="#">8</a>
<a href="#">4.2.2</a>	Message Flow.....	<a href="#">8</a>
<a href="#">4.3</a>	ANCP based Line Configuration.....	<a href="#">9</a>
<a href="#">4.3.1</a>	Goals.....	<a href="#">9</a>
<a href="#">4.3.2</a>	Message Flow.....	<a href="#">10</a>
<a href="#">4.4</a>	ANCP based Transactional Multicast.....	<a href="#">12</a>
<a href="#">4.5</a>	ANCP based OAM.....	<a href="#">13</a>



<a href="#">4.5.1</a>	Message Flow.....	<a href="#">13</a>
5	Access Node Control Protocol (ANCP)	14
<a href="#">5.1</a>	ANCP/TCP connection establishment.....	<a href="#">14</a>
<a href="#">5.2</a>	ANCP Connection keep-alive.....	<a href="#">15</a>
<a href="#">5.3</a>	Capability negotiation.....	<a href="#">16</a>
<a href="#">5.4</a>	GSMP Message Extensions for Access Node Control.....	<a href="#">18</a>
<a href="#">5.4.1</a>	General Extensions.....	<a href="#">18</a>
<a href="#">5.4.2</a>	Topology Discovery Extensions.....	<a href="#">21</a>
<a href="#">5.4.3</a>	Line Configuration Extensions.....	<a href="#">30</a>
<a href="#">5.4.4</a>	OAM Extensions.....	<a href="#">32</a>
<a href="#">5.5</a>	ATM-specific considerations.....	<a href="#">35</a>
<a href="#">5.6</a>	Ethernet-specific considerations.....	<a href="#">36</a>
6	IANA Considerations	36
7	Security Considerations	36
8	References	37
	Author's Addresses	38
	Intellectual Property Statement	38
	Disclaimer of Validity	39
	Copyright Statement	39
	Acknowledgment	39

## 1 Specification Requirements

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

## 2 Introduction

DSL is a widely deployed access technology for Broadband Access for Next Generation Networks. Several specifications like [1-3] describe possible architectures for these access networks. In the scope of these specifications are the delivery of voice, video and data services.

When deploying value-added services across DSL access networks, special attention regarding quality of service and service control is required, which implies a tighter coordination between network elements in the broadband access network without burdening the OSS layer.

This draft defines extensions and modifications to GSMPv3 (specified in [4]) and certain new mechanisms to realize a control plane between a service-oriented layer 3 edge device (the NAS) and a layer2 Access Node (e.g. DSLAM) in order to perform QoS-related, service-related and subscriber-related operations. The control protocol as a result of these extensions and mechanisms is referred to as Access Node Control Protocol (ANCP).

In addition to specifying extensions and modifications to relevant GSMP messages applicable to access node control mechanism, this draft also defines values that ANCP should set for relevant fields in these GSMP messages. However, to understand the basic semantics of various fields in GSMP messages, the reader is referred to [4].

## 3 Broadband Access Aggregation

### 3.1 ATM-based broadband aggregation

End to end DSL network consists of network and application service provider networks (NSP and ASP networks), regional/access network, and customer premises network. Fig1. shows ATM broadband access network components.



The Regional/Access Network consists of the Regional Network, Network Access Server, and the Access Network as show in Fig 1. 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 could consist of DSLAM in the central office, or remote DSLAM, or a Remote Access Multiplexer (RAM). 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 the 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, NAS is also the injection point for policy management and IP QoS in the Regional/Access Networks. In order 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 [1], preserves IP QoS over the ATM network between the NAS and the RGs 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 synch rates.

[9] provides detailed definition and functions of each network element in the broadband reference architecture.

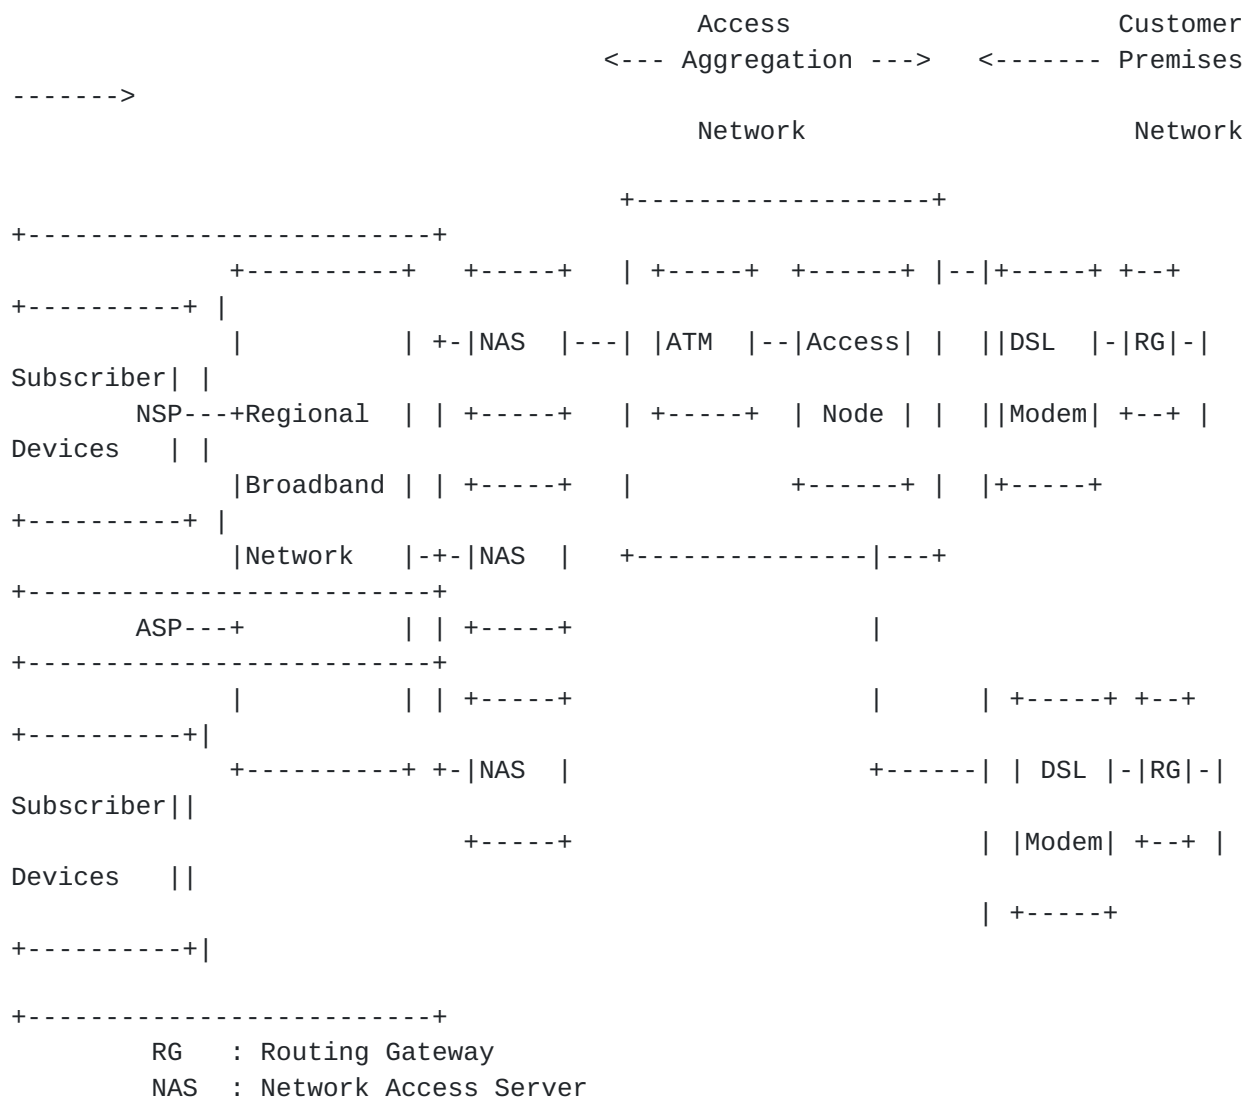


Fig.1 ATM Broadband Aggregation Topology

### 3.2 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 could be used to create a form of virtual path between a given DSLAM and a given NAS. And inner VLAN tags to create a form of virtual circuit on a per DSL line basis. This is 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 N:1 VLAN allocation model. Architectural and topological models of an Ethernet aggregation network in context of DSL aggregation are defined in [8].

## 4 Access Node Control Protocol

### 4.1 Overview

A dedicated control protocol between NAS and access nodes can facilitate "NAS managed" tight QoS control in the access network, simplified OSS infrastructure for service management, optimized multicast replication to enable video services over DSL, subscriber statistics retrieval on the NAS for accounting purposes, and fault isolation capability on the NAS for the underlying access technology. This dedicated control plane is referred to as Access Node Control Protocol (ANCP). This document specifies relevant extensions to GSMPv3 as defined [4] to realize ANCP.

Following sections discuss the use of ANCP for implementing:

- . Dynamic discovery of access topology by the NAS to provide tight QoS control in the access network.
- . Pushing to the access-nodes, subscriber and service data retrieved by the NAS from an OSS system (e.g. radius server), to simplify OSS infrastructure for service management.
- . Optimized, "NAS controlled and managed" multicast replication by access-nodes at L2 layer.
- . NAS controlled, on-demand access-line test capability (rudimentary end-to-end OAM).

In addition to DSL, alternate broadband access technologies (e.g. Metro-Ethernet, Passive Optical Networking, WiMax) will have similar challenges to address, and could benefit from the same approach of a control plane between a NAS and an Access Node (e.g. OLT), providing a unified control and management architecture for multiple access technologies, hence facilitating migration from one to the other and/or parallel deployments.

GSMPv3 is an ideal fit for implementing ANCP. It is extensible and can be run over TCP/IP, which makes it possible to run over different access technologies.



## 4.2 ANCP based Access Topology Discovery

### 4.2.1 Goals

[1] discusses various queuing/scheduling mechanisms to avoid congestion in the access network while dealing with multiple flows with distinct QoS requirements. Such mechanisms require that the NAS gains knowledge about the topology of the access network, the various links being used and their respective rates. Some of the information required is somewhat dynamic in nature (e.g. DSL sync rate), hence cannot come from a provisioning and/or inventory management OSS system. Some of the information varies less frequently (e.g. capacity of a DSLAM uplink), but nevertheless needs to be kept strictly in sync between the actual capacity of the uplink and the image the NAS has of it.

Following section describes ANCP messages that allow the Access Node (e.g. DSLAM) to communicate to the NAS, access network topology information and any corresponding updates.

Some of the parameters that can be communicated from the DSLAM to the NAS include DSL line state, actual upstream and downstream data rates of a synchronized DSL link, maximum attainable upstream and downstream data rates, interleaving delay etc. Topology discovery is specifically important in case the net data rate of the DSL line changes over time. The DSL net data rate may be different every time the DSL modem is turned on. Additionally, during the time the DSL modem is active, data rate changes can occur due to environmental conditions (the DSL line can get "out of sync" and can retrain to a lower value.

### 4.2.2 Message Flow

When a DSL line initially comes up or resynchronizes to a different rate, the DSLAM generates and transmits a GSMP PORT UP EVENT message to the NAS. The extension field in the message carries the TLVs containing DSL line specific parameters. On a loss of signal on the DSL line, a GSMP PORT DOWN message is generated by the DSLAM to the NAS. In order to provide expected service level, NAS needs to learn the initial attributes of the DSL line before the subscriber can log in and access the provisioned services for the subscriber.



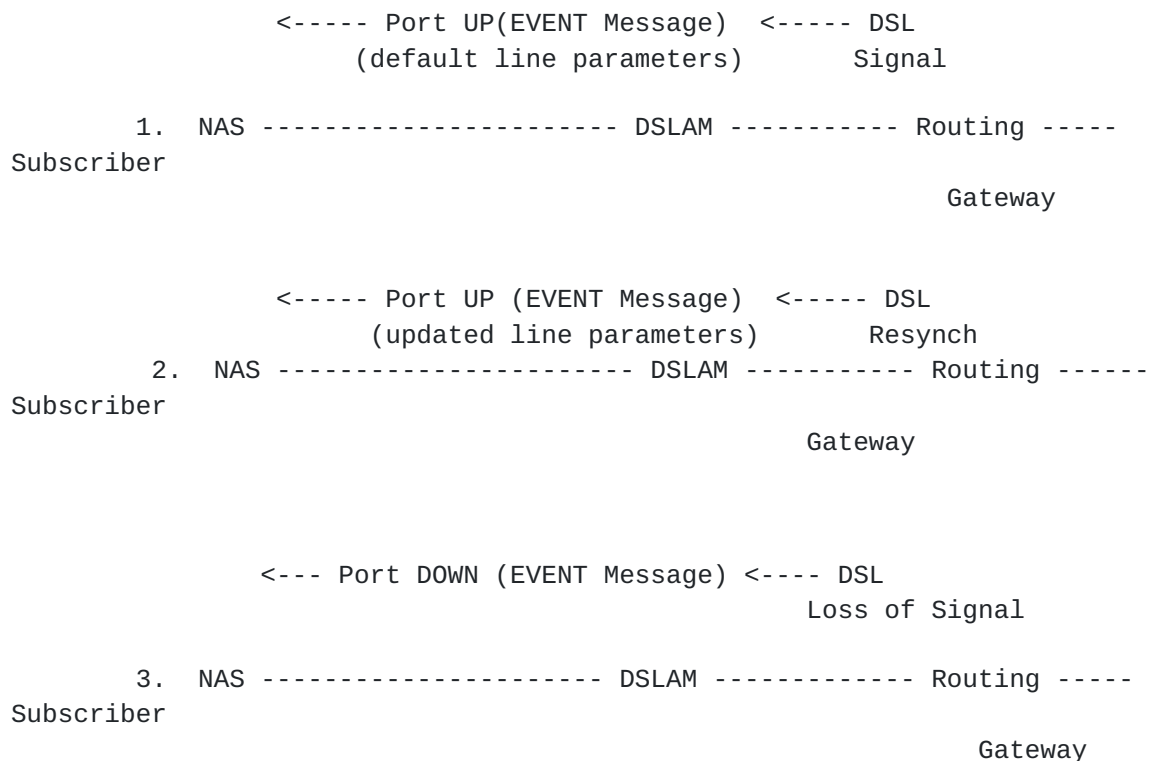


Fig 2. Message flow (ANCP mapping) for topology discovery

The Event message with PORT UP message type (80) is used for conveying DSL line attributes to the NAS. This message with relevant extensions is defined in [section 5.4.2](#).

### 4.3 ANCP based Line Configuration

#### 4.3.1 Goals

Following dynamic discovery of access topology (identification of DSL line and its attributes) as assisted by the mechanism described in the previous section (topology discovery), the NAS could then query a subscriber management OSS system (e.g. RADIUS server) to retrieve subscriber authorization data (service profiles, aka user entitlement). Most of such service mechanisms are typically enforced by the NAS itself, but there are a few cases where it might be useful to push such service parameter to the DSLAM for local enforcement of a mechanism (e.g. DSL-related) on the corresponding subscriber line. One such example of a service parameter that can be pushed to the DSLAM for local enforcement is DSL "interleaving delay". Longer

interleaving delay (and hence stringent error correction) is required



for a video service to ensure better video "quality of experience", whereas for a VoIP service or for "shoot first" gaming service, a very short interleaving delay is more appropriate. Another relevant application is downloading per subscriber multicast channel entitlement information in IPTV applications where the DSLAM is performing IGMP snooping or IGMP proxy function. Using ANCP, the NAS could achieve the goal of pushing line configuration to the DSLAM by an interoperable and standardized protocol.

If a subscriber wants to choose a different service, it can require an OPEX intensive reconfiguration of the line via a network operator, possibly implying a business-to-business transaction between an ISP and an access provider. Using ANCP for line configuration from the NAS dramatically simplifies the OSS infrastructure for service management, allowing fully centralized subscriber-related service data (e.g. RADIUS server back-end) and avoiding complex cross-organization B2B interactions.

The best way to change line parameters would be by using profiles. These profiles (DSL profiles for different services) are pre-configured on the DSLAMs. The NAS can then indicate a reference to the right DSL profile via ANCP. Alternatively, discrete DSL parameters can also be conveyed by the NAS in ANCP.

#### 4.3.2 Message Flow

Triggered by topology information reporting a new DSL line or triggered by a subsequent user session establishment (PPP or DHCP), the NAS may send line configuration information (e.g. reference to a DSL profile) to the DSLAM using GSMP Port Management messages. The NAS may get such line configuration data from a policy server (e.g. RADIUS). Figure 3 summarizes the interaction.

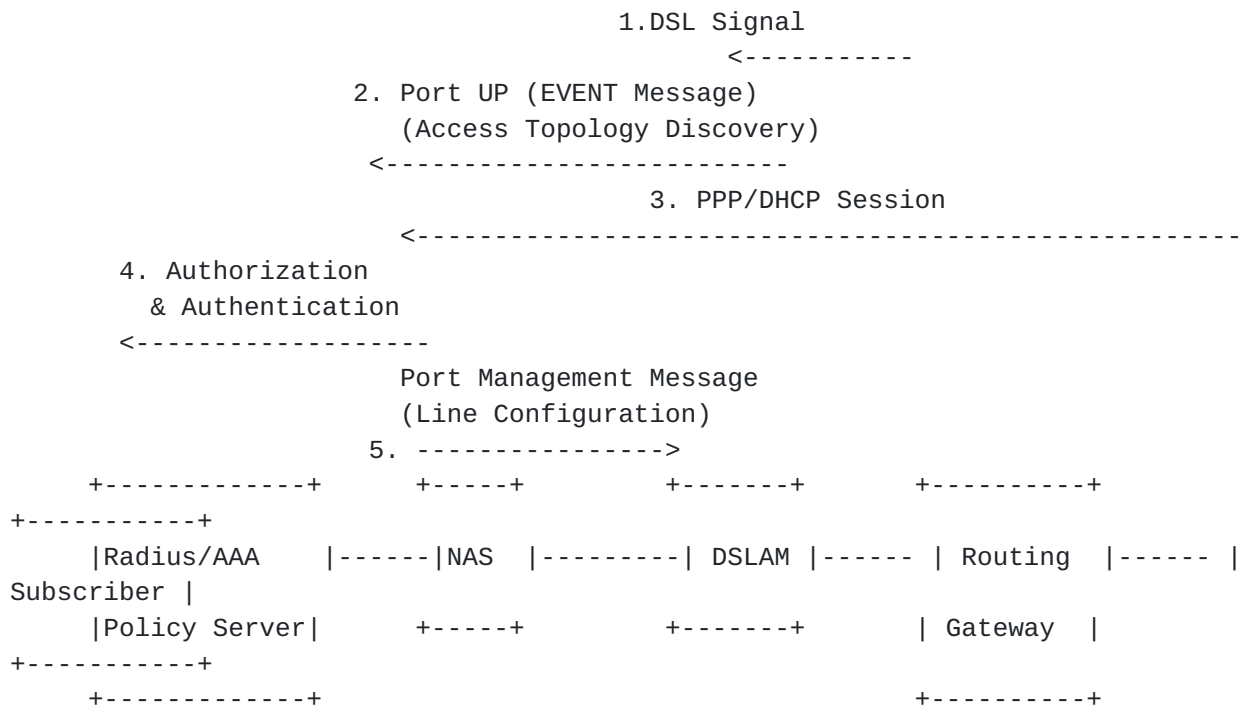


Fig 3. Message flow - ANCP mapping for Initial Line Configuration

The NAS may update the line configuration due to a subscriber service change (e.g. triggered by the policy server). Figure 4 summarizes the interaction



## 1. PPP/DHCP Session

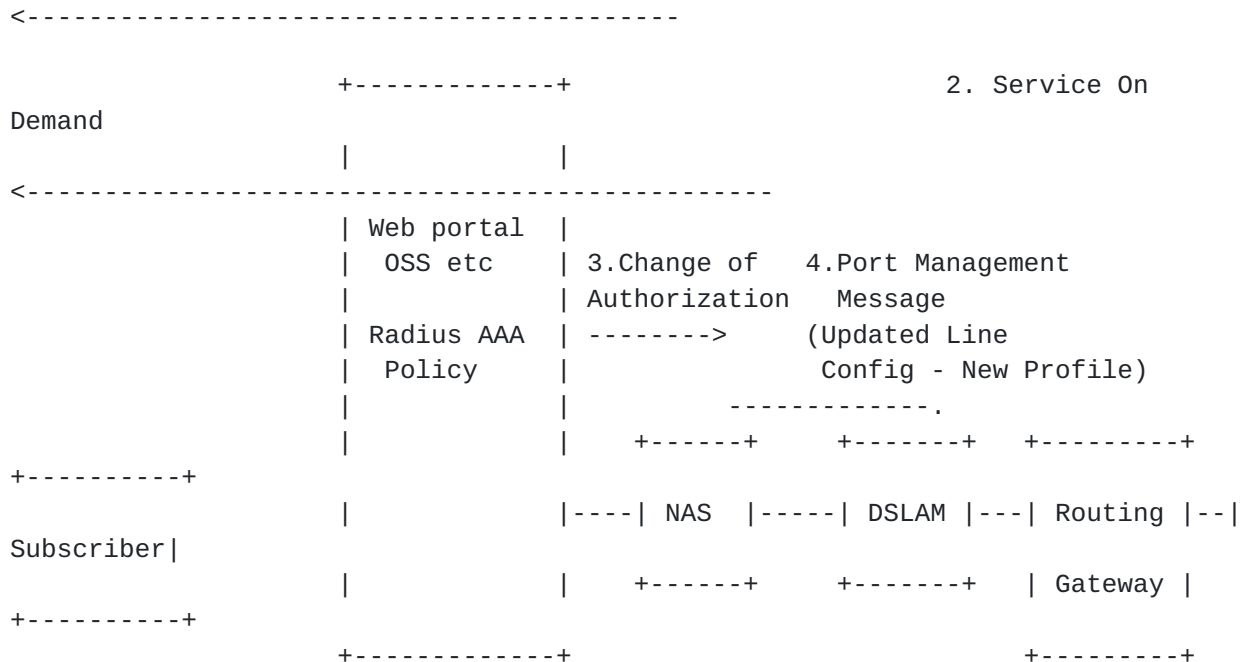


Fig 4. Message flow - ANCP mapping for Updated Line

Configuration

The format of relevant extensions to port management message is defined in [section 5.4.3](#). The line configuration models could be viewed as a form of delegation of authorization from the NAS to the DSLAM.

## 4.4 ANCP based Transactional Multicast

Typical IP multicast in access networks involves the NAS terminating user requests for receiving multicast channels via IGMP. The NAS authorizes the subscriber, and dynamically determines the multicast subscription rights for the subscriber. Based on the user's subscription, the NAS can replicate the same multicast stream to multiple subscribers. This leads to a waste of access bandwidth if multiple subscribers access network services via the same access-node (e.g. DSLAM). The amount of multicast replication is of the order of number of subscribers rather than the number of access-nodes.

It is ideal for NAS to send a single copy of the multicast stream to a given access-node, and let the access-node perform multicast

replication by layer2 means (e.g. ATM point-to-multipoint cell replication or ethernet data-link bridging) for subscribers behind the access-node. However, operationally, NAS is the ideal choice to handle subscriber management functions (authentication, authorization, accounting and address management), multicast policies

such as per-channel authorization, and complex multicast routing protocols. Therefore, some means is needed for the NAS to setup multicast replication state in the access-nodes. In ATM access networks, ANCP can be used by the NAS to setup P2MP cross-connects in the DSLAMs. Protocol support for this use-case will be provided in future.

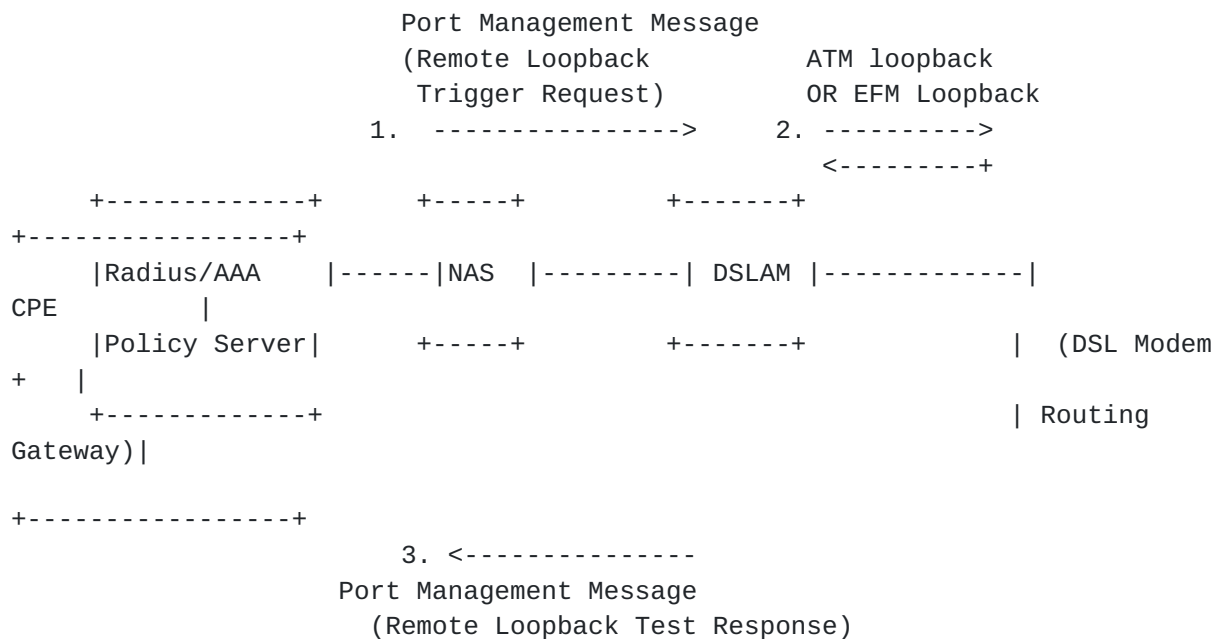
#### 4.5 ANCP based OAM

In a mixed Ethernet and ATM access network (including the local loop), it is desirable to provide similar mechanisms for connectivity checks and fault isolation, as those used in an ATM based architecture. This can be achieved using an ANCP based mechanism until end-to-end Ethernet OAM mechanisms are more widely implemented in various network elements.

A simple solution based on ANCP can provide NAS with an access-line test capability and to some extent fault isolation. Controlled by a local management interface the NAS can use an ANCP operation to trigger the access-node to perform a loopback test on the local-loop (between the access-node and the CPE). The access-node can respond via another ANCP operation the result of the triggered loopback test. In case of ATM based local-loop the ANCP operation can trigger the access-node to generate ATM (F4/F5) loopback cells on the local loop. In case of Ethernet, the access-node can trigger an ethernet loopback message(per EFM OAM) on the local-loop.

##### 4.5.1 Message Flow

Port Management message can be used by the NAS to request access node to trigger a remote loopback test on the local loop. The result of the loopback test can be asynchronously conveyed by the access node to the NAS in a Port Management response message. The format of relevant extensions to port management message is defined in [section 5.4.4](#).



## 5 Access Node Control Protocol (ANCP)

### 5.1 ANCP/TCP connection establishment

ANCP will use TCP for exchanging protocol messages. [5] defines the GSMP message encapsulation for TCP. The TCP session is initiated from the DSLAM (access node) to the NAS (controller). This is necessary to avoid static provisioning on the NAS for all the DSLAMs that are being served by the NAS. It is easier to configure a given DSLAM with the single IP address of the NAS that serves the DSLAM. This is a deviation from [5] which indicates that the controller initiates the TCP connection to the switch.

NAS listens for incoming connections from the access nodes. Port 6068 is used for TCP connection. Adjacency protocol messages, which are used to synchronize the NAS and access-nodes and maintain handshakes, are sent after the TCP connection is established. ANCP messages other than adjacency protocol messages may be sent only after the adjacency protocol has achieved synchronization.

In the case of ATM access, a separate PVC (control channel) capable of transporting IP would be configured between NAS and the DSLAM for ANCP messages.





## 5.2 ANCP Connection keep-alive

GSMPv3 defines an adjacency protocol. The adjacency protocol is used to synchronize states across the link, to negotiate which version of the GSMP protocol to use, to discover the identity of the entity at the other end of a link, and to detect when it changes. GSMP is a hard state protocol. It is therefore important to detect loss of contact between switch and controller, and to detect any change of identity of switch or controller. No protocol messages other than those of the adjacency protocol may be sent across the link until the adjacency protocol has achieved synchronization. There are no changes to the base GSMP adjacency protocol for implementing ANCP.

The NAS will set the M-flag in the SYN message (signifying it is the master). Once the adjacency is established, periodic adjacency messages (type ACK) are exchanged. The default ACK interval as advertised in the adjacency messages is 10 sec for ANCP. It SHOULD be configurable and is an implementation choice. It is recommended that both ends specify the same timer value. However, it is not necessary for the timer values to match.

The GSMP adjacency message defined in [4] is extended for ANCP and is shown in [section 5.3](#) immediately following this section. The 8-bit version field in the adjacency protocol messages is modified to carry the version and sub-version of the GSMP protocol for version negotiation. ANCP uses version 3 and sub-version 1 of GSMP protocol. The semantics and suggested values for Code, Sender Name, Receiver Name, Sender Instance, and Receiver Instance fields are as defined in [4]. The Sender Port, and Receiver Port should be set to 0 by both ends. The pType field should be set to 0. The pFlag should be set to 1.

If the adjacency times out on either end, due to not receiving an adjacency message for a duration of (3 \* Timer value), where the timer value is specified in the adjacency message, all the state received from the ANCP neighbor should be cleaned up, and the TCP connection should be closed. The NAS would continue to listen for new connection requests. The DSLAM will try to re-establish the TCP connection and both sides will attempt to re-establish the adjacency.

The handling defined above will need some modifications when ANCP graceful restart procedures are defined. These procedures will be defined in a separate draft.



### 5.3 Capability negotiation

The adjacency message as defined in [4] is extended to carry technology specific "Capability TLVs". Both the NAS and the access node will advertise supported capabilities in the originated adjacency messages. If a received adjacency message indicates absence of support for a capability that is supported by the receiving device, it will turn off the capability locally and will send an updated adjacency message with the capability turned off to match the received capability set. This process will eventually result in both sides agreeing on the minimal set of supported capabilities. The adjacency will not come up unless the capabilities advertised by the controller and the controlled device match.

After initial synchronization, if at anytime a capability mismatch is detected, the adjacency will be brought down (RSTACK will be generated by the device detecting the mismatch), and synchronization will be re-attempted.

```

      0              1              2              3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Ver | Sub | Message Type | Timer | M | Code |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Sender Name                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Receiver Name                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Sender Port                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Receiver Port                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| PType | PFlag | Sender Instance |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Partition ID | Receiver Instance |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Tech Type | # of TLVs | Total Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     |
~                                     ~
|                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```



The format of capability TLV is:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Capability Type           |   Capability Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
~
~                               Capability Data
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Tech Type field type indicates the technology to which the capability extension applies. For access node control in case of DSL networks, new type "DSL" is proposed. The value for this field is 0x05. This is the first available value in the range that is not currently allocated. It will need to be reserved with IANA.

Capability length is the number of actual bytes contained in the value portion of the TLV. The TLV is padded to a 4-octet alignment. Therefore, a TLV with no data will contain a zero in the length field (if capability data is three octets, the length field will contain a three, but the size of the actual TLV is eight octets).

Capability data field can be empty if the capability is just a boolean. In case the capability is a boolean, it is inferred from the presence of the TLV (with no data). Capability data provides the flexibility to advertise more than mere presence or absence of a capability. Capability types can be registered with IANA. Following capabilities are defined for ANCP as applied to DSL access:

1. Capability Type : Dynamic-Topology-Discovery = 0x01

Length (in bytes) : 0

Capability Data : NULL

2. Capability Type : Line-Configuration = 0x02

Length (in bytes) : 0

Capability Data : NULL

3. Capability Type : Transactional-Multicast = 0x03 (controller i.e. NAS terminates IGMP messages from subscribers, and via l2 control protocol, signals state to the access-nodes (e.g. DSLAMs) to enable layer2 replication of multicast streams. In ATM access network this implies that NAS instructs the access-node to setup a P2MP cross-connect. The details of this will be covered in a separate ID [6].

Length (in bytes) : 0

Capability Data : NULL

4. Capability Type : OAM = 0x04

Length (in bytes) : 0

Capability Data : NULL

5. Capability Type : Bulk Transaction = 0x05 (defined in [section 5.4.1.2](#)).

Length (in bytes) : 0

Capability Data : NULL

## 5.4 GSMP Message Extensions for Access Node Control

### 5.4.1 General Extensions

Extensions to GSMP messages for various use-cases of Access Node Control mechanism are defined in sections [5.4.2](#) to [5.4.4](#). However, sub-sections [5.4.1.1](#) and [5.4.1.2](#) below define extensions to GSMP that have general applicability.

#### 5.4.1.1 Extension TLV

In order to provide flexibility and extensibility certain GSMP messages such as PORT MANAGEMENT and EVENT messages defined in [4] have been modified to include an extension block that follows a TLV structure. Individual messages in the following sections describe the usage and format of the extension block.

All Extension TLV's will be designated as follow:

[illegible]

x: Reserved Flags

These are generally used by specific messages and will be defined in those messages.

Message Type

An 8-bit field corresponding to the message type where the extension block is used.

Tech Type

An 8-bit field indicating the applicable technology type value. The Message Type plus the Tech Value uniquely define a single Extension Type and can be treated as a single 16 bit extension type. Tech Type value of 0x05 SHOULD be used by ANCP for DSL technology.

```
0x00      Extension block not it use.
```

```
0x01    0x04    Already in use by various technologies
```

0x05 DSL

0x06 - 0xFE Reserved







In a Bulk Transaction Message, each of the message in the payload is framed with a complete header and is acted on individually. The response to the Bulk Transaction message contains the response message that would have been generated by each of the messages had it been sent individually. Each response message will have the appropriate result and code field filled. Any message can be included in the bulk Transaction message except for adjacency message and another bulk transaction message. If a prohibited message is included in a bulk Transaction message, it **MUST** be included in the Bulk response with a failure response. The response code for that failure is 0x81 ( Message type prohibited in bulk Transaction ). While the individual message would fail, this would not constitute a failure for the Bulk Transaction message

#### 5.4.2 Topology Discovery Extensions

The GSMP Event message with PORT UP message type (80) is used for conveying DSL line attributes to the NAS. The version field should be set to 3, and the sub field should be set to 1. The I and subMessage fields **SHOULD** be set to 1 to signify no fragmentation. The Port and Label fields should be set to 0. The Port Session Number should be set to 0, and the Event Sequence Number should be 0. The Tech Type field is extended with new type "DSL". The value for this field is 0x05. The message **SHOULD** be generated when a line first comes UP, or any of the attributes of the line change e.g. the line re-trains to a different rate or one or more of the configured line attributes are administratively modified. Also, when the ANCP session first comes up, the DSLSM **SHOULD** transmit a PORT UP message to the NAS for each line that is up. A DSLAM **MAY** use a Bulk Transaction message as defined in 5.4.1.2 to aggregate the transmission of PORT UP messages.

When a DSL line goes down (idle or silent), the DSLAM **SHOULD** transmit an Event message with PORT DOWN message type (81) to the NAS. It is recommended that the DSLAMs use a dampening mechanism per DSL line to control the rate of state changes per DSL line, communicated to the NAS.



```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Vers | Sub | Message Type | Result | Code |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Partition ID | Transaction Identifier |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| I | SubMessage Number | Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Port |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Port Session Number |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Event Sequence Number |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x|S|x|x|
+---+---+---+
Label
~
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|x|x|x|x|x|x|x|x| Message Type | Tech Type | Block Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
~
Extension Value
~
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The format of the "Extension Value" field for Tech Type DSL is as follows :

```


+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| # of TLVs | Extension block length (bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|
~
~
TLVs
~
|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Extension Value contains one or more TLVs to identify a DSL line and define its characteristics. A TLV can consist of multiple sub-TLVs. First 2 byte of the Extension Value contains the number of TLVs that follow. The next 2 bytes contain the total length of the TLVs carried in the extension block in bytes (existing Block Length field in the GSMP message is limited to 255 bytes and is not sufficient).



General format of a TLV is :



The value field in each TLV is padded to a 4-octet alignment. The Length field in each TLV contains the actual number of bytes in the TLV (not including the padding if present). If a TLV is not understood by the NAS, it is silently ignored. Currently defined types start from 0x01.

Following TLVs are currently defined:

1. Type (Access-Loop-Circuit-ID = 0x01): This is a mandatory TLV and contains an identifier of the subscriber's connection to the access node (i.e. local loop). The local loop can be ATM based or Ethernet based. The Access Loop Circuit ID has local significance at the access node. The exact usage on the NAS is beyond the scope of this document. The format used for local loop identification in ANCP messages MUST be identical to what is used by the access nodes in subscriber signaling messages when the access nodes act as signaling relay agents as outlined in [7] and [8].

Length : (upto 63 bytes)

Value : ASCII string.

For an ATM based local loop the string consists of slot/port and VPI/VCI information corresponding to the subscriber's DSL connection. Default syntax for the string inserted by the access node as per [8] is:

Access-Node-Identifier atm slot/port:vpi.vci

The Access-Node-Identifier uniquely identifies the access node in the access network. The slot/port and VPI/VCI uniquely identifies the DSL line on the access node. Also, there is one to one



correspondence between DSL line and the VC between the access node and the NAS.

For local loop which is Ethernet based (and tagged), the string consists of slot/port and VLAN tag corresponding to the subscriber. Default syntax for the string inserted by the access node as per [8] is:

```
"Access-Node-Identifier eth slot/port[:vlan-id]"
```

2. Type (Access-Loop-Remote-Id = 0x02): This is an optional TLV and contains an identifier to uniquely identify a user on a local loop on the access node. The exact usage on the NAS is out of scope of this document. It is desirable that the format used for the field is similar to what is used by the access nodes in subscriber signaling messages when the access nodes act as signaling relay agents as outlined in [7] and [8].

Length : (upto 63 bytes)

Value : ASCII string

3. Type (Access-Aggregation-Circuit-ID-Binary = 0x06)

Length : (8 bytes)

Value : two 32 bit integers.

For ethernet access aggregation, where a per-subscriber (stacked) VLAN can be applied (1:1 model defined in [8]), the VLAN stack provides a convenient way to uniquely identify the DSL line. The outer VLAN is equivalent to virtual path between a DSLAM and the NAS and inner VLAN is equivalent to a virtual circuit on a per DSL line basis. In this scenario, any subscriber data received by the access node and transmitted out the uplink to the aggregation network will be tagged with the VLAN stack assigned by the access node

This TLV can carry the VLAN tags assigned by the access node in the ANCP messages. The VLAN tags can uniquely identify the DSL line being referred to in the ANCP messages, assuming the VLAN tags are not in any way translated in the aggregation network and are unique across physical ports. Each 32 bit integer (least significant bits)





contains a 12 bit VLAN identifier (which is part of the VLAN tag defined by IEEE 802.1Q).

Also, in case of an ATM aggregation network, where the DSLAM is directly connected to the NAS (without an intermediate ATM switch), the two values can contain VPI and VCI on the DSLAM uplink (and can uniquely identify the DSL line on the DSLAM).

This TLV is optional.

#### 4. Type (Access-Aggregation-Circuit-ID-ASCII = 0x03)

Length : (upto 63 bytes)

Value : ASCII string.

This field contains information pertaining to an uplink on the access node. For Ethernet access aggregation, assuming the access node assigns VLAN tags (1:1 model), typical format for the string is:

id] Access-Node-Identifier eth slot/port [:inner-vlan-id][:outer-vlan-

The slot/port corresponds to the ethernet uplink on the access node towards the NAS.

For an ATM aggregation network, typical format for the string is:

Access-Node-Identifier atm slot/port:vpi.vci

This TLV allows the NAS to associate the information contained in the ANCP messages to the DSL line on the access node.

If the access node inserts this string in the ANCP messages, when referring to local loop characteristics (e.g. DSL line in case of a DSLAM), then it should be able to map the information contained in the string uniquely to the local loop (e.g. DSL line).

On the NAS, the information contained in this string can be used to derive an aggregation network facing construct (e.g. an IP interface) corresponding to the local loop (e.g. DSL line). The association could be based on local configuration on the NAS.

The access node can also convey to the NAS, the characteristics (e.g. bandwidth) of the uplink on the access node. This TLV then



serves the purpose of uniquely identifying the uplink whose characteristics are being defined. A separate set of sub-TLVs will be defined for the uplink characteristics (TBD).

This TLV is optional.

5. Type (DSL Line Attributes = 0x04):

Length : variable (up to 1024 bytes)

Value : This consists of one or more Sub-TLVs corresponding to DSL line attributes. No sub-TLVs other than the DSL type and DSL line state SHOULD be included in a PORT DOWN message.

The general format of the sub-TLVs is identical to the general TLV format. The value field in each sub-TLV is padded to a 4-octet alignment. The Length field in each sub-TLV contains the actual number of bytes in the TLV (not including the padding if present). Current defined sub-TLV types are start from 0x81.

Following sub-TLVs are currently defined :

- . Type (DSL-Type = 0x91) : Defines the type of transmission system in use. This is a mandatory TLV.

Length : (4 bytes)

Value : (Transmission system : ADSL1 = 0x01, ADSL2 = 0x02, ADSL2+ = 0x03, VDSL1 = 0x04, VDSL2 = 0x05, SDSL = 0x06, UNKNOWN = 0x07).

- . Type (Actual-Data-Rate-Upstream = 0x81) : Actual data rate upstream of a synchronized DSL line. This is a mandatory TLV. This rate value and all the subsequent ones account for the DSL overhead (i.e. signify the net rate).

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Actual-Data-Rate-Downstream = 0x82) : Actual data rate downstream of a synchronized DSL line. This is a mandatory TLV.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Minimum-Data-Rate-Upstream = 0x83) : Minimum data rate desired by the operator. This is optional.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Minimum-Data-Rate-Downstream = 0x84) : Minimum data rate desired by the operator. This is optional.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Attainable-Data-Rate-Upstream = 0x85) : Maximum upstream rate that can be attained on the DSL line. This is an optional TLV.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Attainable-Data-Rate-Downstream = 0x86) : Maximum downstream rate that can be attained on the DSL line. This is an optional TLV.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Maximum-Data-Rate-Upstream = 0x87) : Maximum data rate desired by the operator. This is optional.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Maximum-Data-Rate-Downstream = 0x88) : Maximum data rate desired by the operator. This is optional.

Length : (4 bytes)

Value : (Rate in Kb/sec)

- . Type (Minimum-Low-Power-Data-Rate-Upstream = 0x89) : Minimum data rate desired by the operator in low power state. This is optional.  
  
Length : (4 bytes)  
  
Value : (Rate in Kb/sec)
- . Type (Minimum-Low-Power-Data-Rate-Downstream = 0x8A) : Minimum data rate desired by the operator in low power state. This is optional.  
  
Length : (4 bytes)  
  
Value : (Rate in Kb/sec)
- . Type (Maximum-Interleaving-Delay-Upstream = 0x8B) : maximum one way interleaving delay. This is optional.  
  
Length : (4 bytes)  
  
Value : (Time in msec)
- . Type (Actual-Interleaving-Delay-Upstream = 0x8C) : Value corresponding to the interleaver setting. This is optional.  
  
Length : (4 bytes)  
  
Value : (Time in msec)
- . Type (Maximum-Interleaving-Delay-Downstream = 0x8D) : maximum one way interleaving delay. This is optional.  
  
Length : (4 bytes)  
  
Value : (Time in msec)
- . Type (Actual-Interleaving-Delay-Downstream = 0x8E) : Value corresponding to the interleaver setting. This is optional.  
  
Length : (4 bytes)  
  
Value : (Time in msec)

- . Type (DSL line state = 0x8F) : The state of the DSL line. For PORT UP message, at this time, the TLV is optional (since the message type implicitly conveys the state of the line). For PORT DOWN, the TLV is mandatory, since it further communicates the state of the line as IDLE or SILENT.

Length : (4 bytes)

Value : { SHOWTIME = 0x01, IDLE = 0x02, SILENT = 0x03 }

- . Type (Access Loop Encapsulation = 0x90) : The data link protocol and, optionally the encapsulation overhead on the access loop. This is an optional TLV. However, when this TLV is present, the data link protocol MUST minimally be indicated. The encapsulation overhead can be optionally indicated.

Length : (3 bytes)

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

Data Link (1 byte): {ATM AAL5 = 0,  
ETHERNET = 1}

Encaps 1 (1 byte): {NA = 0,  
Untagged Ethernet = 1,  
Single-tagged Ethernet = 2}

Encaps 2 (1 byte):{ 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,







The format of the "Extension Value" field is as follows:

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           # of TLVs           | Extension Block length (bytes) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                                           |
~                                                                           ~
~                               TLVs                                     ~
|                                                                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The Extension Value field contains one or more TLVs containing DSL line identifier and desired service attributes of the the DSL line. First 2 byte of the Extension Value contains the number of TLVs that follow. The next 2 bytes contain the total length of the extension block in bytes (existing Block Length field in the GSMP message is limited to 255 bytes and is not sufficient).

General format of a TLV is:

```

 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type           |           Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                                           |
~                                                                           ~
~                               Value                                     ~
|                                                                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

The value field is padded to a 4-octet alignment. The Length field in each TLV contains the actual number of bytes in the TLV (not including the padding if present). If a TLV is not understood by the access-node, it is silently ignored. Depending upon the deployment scenario, the NAS may specify Access Loop Circuit-ID or the Access Aggregation Circuit-ID) as defined in [section 5.4.1](#). Following TLVs can appear in this message:

- o Type (Access-Loop-Circuit-ID = 0x01) : defined in [section 5.4.1](#)



- o Type (Access-Aggregation-Circuit-ID-Binary = 0x06): defined in [section 5.4.1](#).
- o Type (Access-Aggregation-Circuit-ID-ASCII = 0x03): defined in [section 5.4.1](#).
- o Type (Service-Profile-Name = 0x05): Reference to a pre-configured profile on the DSLAM that contains service specific data for the subscriber.

Length : (upto 64 bytes)

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

In future, more TLVs MAY be defined for individual service attributes of a DSL line (e.g. rates, interleaving delay, multicast channel entitlement access-list etc).

#### 5.4.4 OAM Extensions

GSMP Port Management message (type 32) SHOULD be used by the NAS to trigger access node to run a loopback test on the local loop. The message format is defined in [section 5.4.2](#). The version field SHOULD be set to 3 and sub-version field SHOULD be set to 1. The remaining fields in the GSMP header have standard semantics. The function type used in the request message SHOULD be set to remote loopback (type = 0x09). The port, port session number, event sequence number, duration, event flags, flow control flags and code fields SHOULD all be set to 0. The result field SHOULD be set to AckAll to indicate requirement for the access node to send a success for failure response. The transaction ID SHOULD contain a sequence number inserted by the NAS in each request that it generates. The extension field format is also defined above in [section 5.4.2](#). The extension value field can contain one or more TLVs including the access-line identifier on the DSLAM and OAM test characteristics desired by the NAS.

The TLV format is defined above in [section 5.4.2](#). The value field is padded to a 4-octet alignment. The Length field in each TLV contains the actual number of bytes in the TLV (not including the padding if present). If a TLV is not understood by the NAS, it is silently ignored. Depending upon the deployment scenario, the NAS may specify Access Loop Circuit-ID or the Access Aggregation Circuit-ID) as defined in [section 5.4.1](#). Following TLVs can appear in this message:



- o Type (Access-Loop-Circuit-ID = 0x01) : defined in [section 5.4.1](#)
- o Type (Access-Aggregation-Circuit-ID-Binary = 0x06): defined in [section 5.4.1](#).
- o Type (Access-Aggregation-Circuit-ID-ASCII = 0x03): defined in [section 5.4.1](#).
- o Type (OAM-Loopback-Test-Parameters = 0x07): Parameters related to loopback test. This is an optional TLV. If this TLV is not present in the request message, the DSLAM SHOULD use locally determined default values for the test parameters.

Length: 4 bytes

Value : two 1 byte numbers described below (listed in order of most to least significant). Thus, the 4 bytes consist of 1 byte of Count, followed by 1 byte of Timeout, followed by two pad bytes of zero.

- o Count (1 byte) : Number of loopback cells/messages that should be generated on the local loop as part of the loopback test. The NAS SHOULD restrict the count to be greater than 0 and less than or equal to 32. The DSLAM SHOULD discard the request for a loopback test, if the received test parameters contain an out of range value for the count field. The DSLAM MAY optionally send a failure response to the NAS with the code invalid test parameter .
  - o Timeout (1 byte) : Upper bound on the time in seconds that the NAS would wait for a response from the DSLAM. If the total time taken by the DSLAM to complete a test with requested parameters, exceeds the specified timeout value, it can choose to omit the generation of a response to the NAS. DSLAM SHOULD use a locally determined value for the timeout , if the received value of the timeout parameter is 0.
- 
- o Type (Opaque-Data = 0x08) : This is an optional TLV. If present in the request message, the DSLAM SHOULD reflect it back in the response unmodified.

Length : 8 bytes



Value : Two 32 bit integers inserted by the NAS (not to be interpreted by the DSLAM, but just reflected back in the response).

The access node generates a success or failure response when it deems the loopback test to be complete. Port Management message (type 32) is used. The result field SHOULD be set to success or failure. The function type SHOULD be set to 0x09. The transaction ID SHOULD be copied from the sequence number contained in the corresponding request. The other parameters not explicitly defined here SHOULD be set as specified in the request message above. The code field SHOULD be set to a value in the range 0x500 to 0x5ff (to be reserved with IANA) to indicate the status of the executed test. The valid values defined are (can be extended in future):

0x500 : Specified access line does not exist.

0x501 : Loopback test timed out.

0x502 : Reserved

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.

0x509 : Invalid test parameter.

The Extension value can contain one or more TLVs including the TLV to identify the access line on which the test was performed, and details from executing the test. The access line identifier SHOULD be identical to what was contained in the request. The relevant TLVs are:

- o Type (Access-Loop-Circuit-ID = 0x01) : defined in [section 5.4.1](#)
- o Type (Access-Aggregation-Circuit-ID-Binary = 0x06): defined in [section 5.4.1](#).
- o Type (Access-Aggregation-Circuit-ID-ASCII = 0x03): defined in [section 5.4.1](#).





- o Type (Opaque-Data = 0x08) : Data inserted by the NAS in the request reflected back by the DSLAM.

Length : 8 bytes

Value : Two 32 bit integers as received in the request (opaque to the DSLAM).

- o Type (OAM-Loopback-Test-Response-String = 0x09)

Length (upto 128 bytes)

Value : Suitably formatted ASCII string containing useful details about the test that the NAS will display for the operator, exactly as received from the DSLAM (no manipulation/interpretation by the NAS). This is an optional TLV, but it is strongly recommended, that in case of ATM based local loop, the DSLAM at the very least indicates via this TLV, the total loopback cells generated and the total loopback cells successfully received as part of executing the requested loopback test.

## 5.5 ATM-specific considerations

The topology discovery and line configuration involve the DSL line attributes. For ATM based access networks, the DSL line on the DSLAM is identified by the port and PVP/PVC corresponding to the subscriber. The DSLAMs are connected to the NAS via an ATM access aggregation network. Since, the DSLAM (access-node) is not directly connected to the NAS, the NAS needs a mechanism to learn the DSL line identifier (more generally referred to as "Access Loop Circuit-ID") corresponding to a subscriber. The "Access loop circuit-ID" has no local significance on the NAS. The ANCP messages for topology discovery and line configuration carry opaque "Access loop Circuit-ID" which has only local significance on the DSLAMs.

The access loop circuit identifier can be carried as an ASCII string in the ANCP messages. This allows ANCP to be decoupled from the specifics of the underlying access technology being controlled. On the other hand, this requires a NAS mechanism by which such identifier can be correlated to the context of an aggregation network facing IP interface (corresponding to the subscriber) on the NAS. This would typically require local configuration of such IP interfaces, or of the underlying ATM interfaces.



## 5.6 Ethernet-specific considerations

One possible way of approaching the use of Ethernet technology in the access aggregation network is to recreate the equivalent of Virtual Paths (VPs) and Virtual Circuits (VCs) by using stacked Virtual LAN tags. As an example, one could use an outer VLAN to create a form of virtual path between a given DSLAM and a given NAS. And then use inner VLAN tags to create a form of virtual circuit on a per DSL line basis. In this case, VLAN tags conveyed in topology discovery and line configuration messages will allow to uniquely identify the DSL line in a straightforward manner, assuming the VLAN tags are not translated in some way by the aggregation network, and are unique across physical ports.

However, some carriers do not wish to use this connection oriented approach. Therefore, an alternative model is to bridge sessions from multiple subscribers behind a DSLAM to a single VLAN in the aggregation network. This is the N:1 model. In this model, or in the case where user traffic is sent untagged, the access node needs to insert the exact identity of the DSL line in the topology discovery and line configuration messages, and then have a mechanism by which this can be correlated to the context of an aggregation network facing IP interface (for the subscriber) on the NAS. This can either be based on local configuration on the NAS, or on the fact that such DSLAM (access node) typically inserts the Access Loop Circuit ID in subscriber signaling messages relayed to the NAS (i.e. DHCP or PPPoE discovery messages).

[Section 5.4.1](#) defines Access Loop Circuit ID .

## 6 IANA Considerations

New Tech-Type, capability types, sub-TLV types related to topology discovery and line configuration will need to be reserved.

## 7 Security Considerations

The NAS and DSLAMs are implicitly in a trusted domain, so security for ANCP is not a strong requirement. However, if needed security can be provided using IP security as indicated in [[RFC3293](#)].



## 8 References

- [1] DSLForum TR-059, DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services, Tom Anschutz (BellSouth Telecommunications), 09/2003.
- [2] DSLForum TR-058, Multi-Service Architecture & Framework Requirements, Mark Elias (SBC) and Sven Ooghe (Alcatel), 09/2003.
- [3] DSLForum TR-092, Broadband Remote access server requirements document.
- [4] Doria, A. et al, "General Switch Management Protocol- V3" (GSMP v3), [RFC 3292](#), June 2002.
- [5] 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.
- [6] GSMP extensions for ANCP - Transactional Multicast, [draft-moisand-gsmp-ancp-multicast-00](#) (work in progress).
- [7] DHCP Relay Agent Information Option , [RFC 3046](#), January 2001.
- [8] Architecture & Transport: Migration to Ethernet Based DSL Aggregation, DSL Forum WT-101, Cohen et al.
- [9] Framework for Access Node Control Mechanism in Broadband Networks, [draft-ietf-ancp-framework-00.txt](#).

#### Author's Addresses

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

Email: [swadhwa@juniper.net](mailto:swadhwa@juniper.net)

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

Email: [jmoisand@juniper.net](mailto:jmoisand@juniper.net)

Swami Subramanian  
Juniper Networks  
10 Technology Park Drive  
Westford, MA 01886

Email: [ssubramanian@juniper.net](mailto:ssubramanian@juniper.net)

Thomas Haag  
T-systems

Email: [thomas.haag@t-systems.com](mailto:thomas.haag@t-systems.com)

Norber Voigt  
Siemens

Email: [norbert.voigt@siemens.com](mailto:norbert.voigt@siemens.com)

#### Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF

THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgment

Thanks to Peter Arberg, Josef Froehler, Derek Harkness, Kim Hyldgaard, Sandy Ng, Robert Peschi, and Michel Platnic for their input to this document.