Network Working Group   S. Ooghe
Internet-Draft   Alcatel-Lucent
Intended status: Informational   N. Voigt
Expires: August 11, 2010   Nokia Siemens Networks
                                                           M. Platnic
                                                          ECI Telecom
                                                             T. Haag
                                                     Deutsche Telekom
                                                           S. Wadhwa
                                                     Juniper Networks
                                                     February 7, 2010

        **Framework and Requirements for an Access Node Control Mechanism in
                    Broadband Multi-Service Networks
                    draft-ietf-ancp-framework-13.txt**

Abstract

   The purpose of this document is to define a framework 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. The Access Node Control Mechanism will ensure that the
   transmission of the information does not need to go through distinct
   element managers but rather using a direct device-device
   communication. This allows for performing access link related
   operations within those network elements, while avoiding impact on
   the existing Operational Support Systems (OSSes).

   This document first identifies a number of use cases for which the
   Access Node Control Mechanism may be appropriate. It then presents
   the requirements for the Access Node Control Protocol (ANCP) that
   must be taken into account during protocol design. Finally, it
   describes requirements for the network elements that need to support
   ANCP and the described use cases. These requirements should be seen
   as guidelines rather than as absolute requirements. RFC 2119
   therefore does not apply to the nodal requirements.

Status of this Memo

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

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

Drafts.

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

The list of current Internet-Drafts can be accessed at
http://www.ietf.org/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 11, 2010.

Copyright Notice

Table of Contents

1. Introduction

   Digital Subscriber Line (DSL) technology is widely deployed for
   Broadband Access for Next Generation Networks. Several documents
   like Broadband Forum TR-058 [TR-058], Broadband Forum TR-059 [TR-059]
   and Broadband Forum TR-101 [TR-101] describe possible architectures
   for these access networks. The scope of these specifications
   consists of the delivery of voice, video and data services. The
   framework defined by this document is targeted at DSL-based access
   (either by means of ATM/DSL or as Ethernet/DSL). The framework shall
   be open to other access technologies, such as Passive Optical
   Networks using DSL technology at the Optical Network Unit (ONU), or
   wireless technologies like IEEE 802.16. Several Use Cases such as
   Access Topology Discovery, Remote Connectivity Test and Multicast may
   be applied to these access technologies, but the details of this are
   outside the scope of this document.

   Traditional architectures require Permanent Virtual Circuit(s) per
   Subscriber. Such virtual circuit is configured on layer 2 and
   terminated at the first layer 3 device (e.g. Broadband Remote Access
   Server (BRAS)). Beside the data plane, the models define the
   architectures for element, network and service management.
   Interworking at the management plane is not always possible because
   of the organizational boundaries between departments operating the
   local loop, departments operating the ATM network and departments
   operating the IP network. Besides, management networks are usually
   not designed to transmit management data between the different
   entities in real time.

   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 Nodes
   (e.g. Access Nodes and Network Access Server (NAS)), without
   burdening the Operational Support System (OSS) with unpractical
   expectations.

   Therefore, there is a need for an Access Node Control Mechanism
   between a 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. The Access Node Control Mechanism will ensure
   that the transmission of the information does not need to go through
   distinct element managers but rather using a direct device-device
   communication. This allows for performing access link related
   operations within those network elements, while avoiding impact on
   the existing OSS systems.

   This document provides a framework for such an Access Node Control

Mechanism and identifies a number of use cases for which this
mechanism can be justified. Next, it presents a number of
requirements for the Access Node Control Protocol (ANCP) and the
network elements that need to support it.

The requirements spelled out in this document are based on the work
that is performed by the Broadband Forum ([TR-147]).

1.1. Requirements Notation

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].

1.2. Definitions

o Access Node (AN): Network device, usually located at a service
   provider central office or street cabinet, that terminates Access
   Loop connections from Subscribers. In case the Access Loop is a
   Digital Subscriber Line (DSL), this is often referred to as a DSL
   Access Multiplexer (DSLAM).

o Network Access Server (NAS): Network device which aggregates
   multiplexed Subscriber traffic from a number of Access Nodes. The
   NAS plays a central role in per-subscriber policy enforcement and
   QoS. Often referred to as a Broadband Network Gateway (BNG) or
   Broadband Remote Access Server (BRAS). A detailed definition of
   the NAS is given in [RFC2881].

o "Net Data Rate": 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 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.

o "Line Rate": defined by ITU-T G.993.2. It contains the complete
   overhead including Reed-Solomon and Trellis coding.

o Access Node Control Mechanism: a method for multiple network
   scenarios with an extensible communication scheme that conveys
   status and control information between one or more ANs and one or
   more NASs without using intermediate element managers.

o Control Channel: a bidirectional IP communication interface
   between the controller function (in the NAS) and the reporting/

enforcement function (in the AN). It is assumed that this
interface is configured (rather than discovered) on the AN and the
NAS.

o Access Node Control Adjacency: the relationship between an Access
   Node and a NAS 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.

o Multicast Flow: designates datagrams sent to a group from a set of
   sources for which multicast reception is desired. A distinction
   can be made between Any Source Multicast (ASM) and Source Specific
   Multicast (SSM).

o Join: signaling from the user equipment that it wishes to start
   receiving a new multicast flow. In ASM, it is referred to as a
   "Join". In SSM [RFC4607], it is referred to as a "subscribe". In
   IGMPv2, "joins" are indicated through an "IGMPv2 membership
   report". In IGMPv3 [RFC3376], "join" are indicated through
   "membership report" using different Filter-Mode-Change (ASM) and
   Source-List-Change Records.

o Leave: signaling from the user equipment that it wishes to stop
   receiving a multicast flow. With IGMPv2 this is conveyed inside
   the "Leave Group" message while in IGMPv3, "leave" is indicated
   through "IGMPv3 membership report" message using different Filter-
   Mode-Change (ASM) and Source-List-Change Records.

2. General Architecture Aspects

   This section introduces the basic concept of the Access Node Control
   Mechanism and describes the reference architecture where it is being
   applied. Based on the reference architecture, the section then
   describes how Access Node Control messages are to be prioritized over
   other data traffic, and the interaction between ANCP and the network
   management system. Finally, the addressing schemes are described
   that allow identifying an Access Port in Access Node Control
   messages.

2.1. Concept of an Access Node Control Mechanism

   The high-level communication framework for an Access Node Control
   Mechanism is defined in Figure 1. The Access Node Control Mechanism
   defines a quasi-realtime, general-purpose method for multiple network
   scenarios with an extensible communication scheme, addressing the
   different use cases that are described throughout this document.
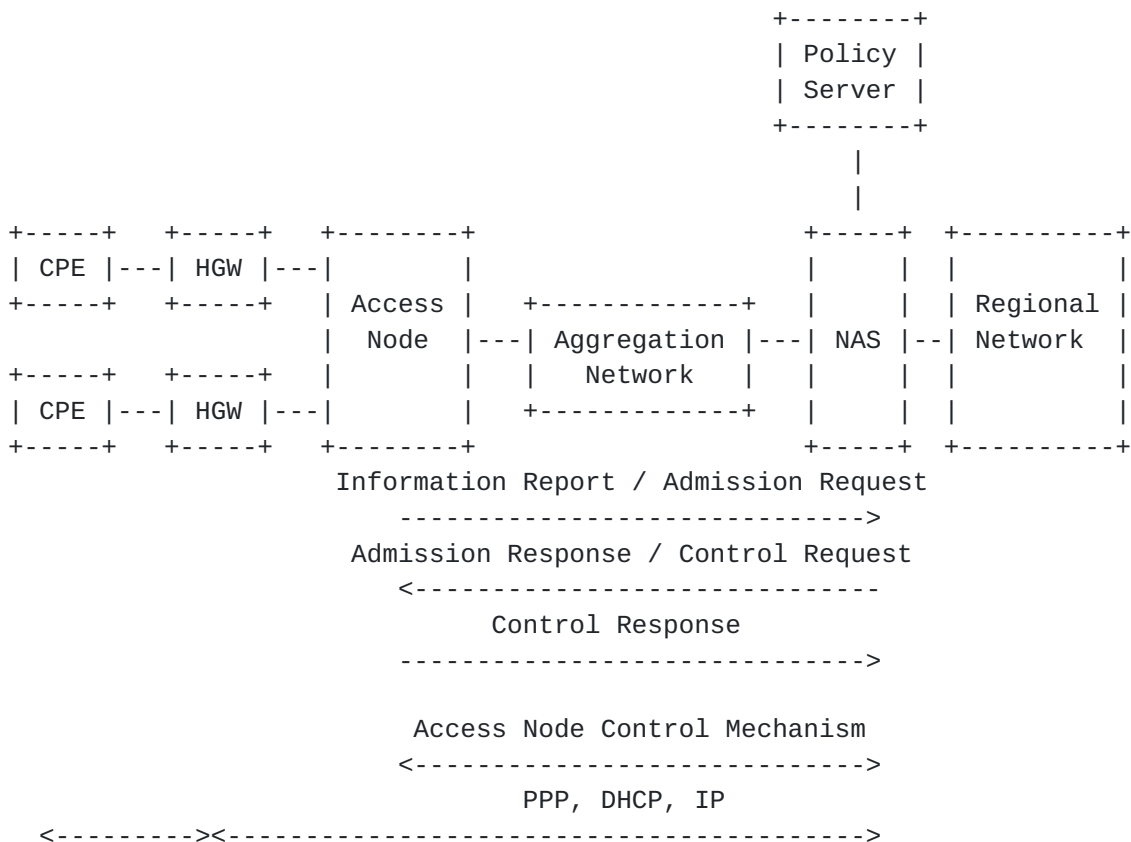
```
                                          +--------+
                                          | Policy |
                                          | Server |
                                          +--------+
                                              |
                                              |
+-----+   +-----+   +--------+             +-----+  +----------+
| CPE |---| HGW |---|        |             |     |  |          |
+-----+   +-----+   | Access |   +-------------+  |     |  | Regional |
                    | Node   |---| Aggregation |---| NAS |--| Network  |
+-----+   +-----+   |        |   |  Network    |  |     |  |          |
| CPE |---| HGW |---|        |   +-------------+  |     |  |          |
+-----+   +-----+   +--------+             +-----+  +----------+
                     Information Report / Admission Request
                     ------------------------------->
                     Admission Response / Control Request
                     <------------------------------
                           Control Response
                     ------------------------------>


                     Access Node Control Mechanism
                     <------------------------------>
                           PPP, DHCP, IP
   <--------><-------------------------------------------->

CPE: Customer Premises Equipment
HWG: Home Gateway
```

                    Figure 1: Access Network Architecture

A number of functions can be identified:

o A controller function: this function is used to either send out
   requests for information to be used by the network element where
   the controller function resides, or to trigger a certain behavior
   in the network element where the reporting and/or enforcement
   function resides.

o A reporting function: this function is used to convey status
   information to the controller function. An example of this is the
   transmission of the Access Loop data rate from an Access Node to a
   Network Access Server (NAS) tasked with shaping traffic to that
   rate.

o An enforcement function: this function is contacted by the
   controller function to trigger a remote action on the Access Node.
   An example is the initiation of a port testing mechanism on an
   Access Node. Another example is enforcing whether a multicast
   join is to be honored or denied.

The messages shown in Figure 1 show the conceptual message flow. The
actual use of these flows, and the times or frequencies when these
messages are generated depends on the actual use case, which are
described in Section 3.

The use cases in this document are described in an abstract way,
independent from any actual protocol mapping. The actual protocol
specification is out of scope of this document, but there are certain
characteristics of the protocol required such as to simplify
specification, implementation, debugging & troubleshooting, but also
to be easily extensible in order to support additional use cases.

2.2. Reference Architecture

The reference architecture used in this document can be based on ATM
or Ethernet access/aggregation. Specifically:

o In case of a legacy ATM aggregation network that is to be used for
   the introduction of new QoS-enabled IP services, the architecture
   builds on the reference architecture specified in Broadband Forum
   [TR-059];

o In case of an Ethernet aggregation network that supports new QoS-
   enabled IP services (including Ethernet multicast replication),
   the architecture builds on the reference architecture specified in
   Broadband Forum [TR-101].

Given the industry's move towards Ethernet as the new access and

aggregation technology for triple play services, the primary focus
throughout this document is on a TR-101 architecture. However the
concepts are equally applicable to an ATM architecture based on TR-
059.

2.2.1. Home Gateway

The Home Gateway (HGW) connects the different Customer Premises
Equipment (CPE) to the Access Node and the access network. In case
of DSL, the HGW is a DSL Network Termination (NT) that could either
operate as a layer 2 bridge or as a layer 3 router. In the latter
case, such a device is also referred to as a Routing Gateway (RG).

2.2.2. Access Loop

The Access Loop ensures physical connectivity between the HGW at the
customer premises, and the Access Node. In case of DSL, the Access
Loop physical layer could be e.g. ADSL, ADSL2+, VDSL, VDSL2 or
SHDSL. In order to increase bandwidth, it is also possible that
multiple DSL links are grouped together to form a single virtual
link; this process is called "DSL bonding". The protocol
encapsulation on the Access Loop could be based on multi-protocol
encapsulation over AAL5, defined in [RFC2684]. This covers PPP over
Ethernet (PPPoE, defined in [RFC2516]), bridged IP (IPoE, defined in
[RFC894]) and routed IP (IPoA, defined in [RFC2225]). Next to this,
PPPoA as defined in [RFC2364] can be used. Future scenarios include
cases where the Access Loop supports direct Ethernet encapsulation
(e.g. when using VDSL or VDSL2).

2.2.3. Access Node

The Access Node (AN) may support one or more Access Loop technologies
and allow them to inter-work with a common aggregation network
technology. Besides the Access Loop termination the AN can also
aggregate traffic from other Access Nodes using ATM or Ethernet.

The framework defined by this document is targeted at DSL-based
access (either by means of ATM/DSL or as Ethernet/DSL). The
framework shall be open to non-DSL technologies, like Passive Optical
Networks (PON) and IEEE 802.16 (WiMAX), but the details of this are
outside the scope of this document.

The reporting and/or enforcement function defined in Section 2.1
typically resides in an Access Node.

2.2.4. Access Node Uplink

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

2.2.5. Aggregation Network

The aggregation network provides traffic aggregation towards the NAS. The aggregation technology can be based on ATM (in case of a TR-059 architecture) or Ethernet (in case of a TR-101 architecture).

2.2.6. Network Access Server

The Network Access Server (NAS) interfaces to the aggregation network by means of standard ATM or Ethernet interfaces, and towards the Regional Network by means of transport interfaces for Ethernet frames (e.g. GigE, Ethernet over SONET). The NAS functionality correpsonds to the BNG functionality described in Broadband Forum TR-101. In addition to this, the NAS supports the Access Node Control functionality defined for the respective use cases throughout this document.

The controller function defined in Section 2.1 typically resides in a NAS.

2.2.7. Regional Network

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

The Regional Network supports aggregation of traffic from multiple Access Networks and hands off larger geographic locations to NSPs and ASPs - relieving a potential requirement for them to build infrastructure to attach more directly to the various Access Networks.

2.3. Prioritizing Access Node Control traffic

When sending Access Node Control messages across the aggregation network, care is needed that messages won't get lost. The connectivity between the Access Node and the NAS may differ depending

on the actual layer 2 technology used (ATM or Ethernet). This
section briefly outlines how network connectivity can be established.

In case of an ATM access/aggregation network, a typical practice is
to send the Access Node Control Protocol messages over a dedicated
Permanent Virtual Circuit (PVC) configured between the AN and the
NAS. These ATM PVCs would then be given a high priority so that at
times of network congestion, loss of the ATM cells carrying the
Access Node Control Protocol is avoided or minimized. It is
discouraged to route the Access Node Control Protocol messages within
the Virtual Path (VP) that also carries the customer connections, if
that VP is configured with a best effort QoS class (e.g. Unspecified
Bitrate (UBR)). The PVCs of multiple Access Node Control Adjacencies
can be aggregated into a VP that is given a high priority and runs
across the aggregation network. This requires the presence of a VC
cross-connect in the aggregation node that terminates the VP.

In 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). This can be achieved using a different VLAN ID for each Access
Node, or, in networks with many Access Nodes and high degree of
aggregation, one Customer VLAN (C-VLAN) per Access Node and one
Service VLAN (S-VLAN) for the Access Node Control Adjacencies of all
Access Nodes. The traffic should be given a high priority (e.g. by
using a high Class of Service (COS) value) so that the frame loss of
Ethernet frames carrying the Access Node Control Protocol messages is
minimized in the event of network congestion.

In both cases, the Control Channel between NAS and Access Node could
use the same physical network- and routing resources as the
Subscriber traffic. This means that the connection is an inband
connection between the involved network elements. Therefore there is
no need for an additional physical interface to establish the Control
Channel.

Note that these methods for transporting Access Node Control Protocol
messages are typical examples; they do not rule out other methods
that achieve the same behavior.

The Access Node Control Adjacency interactions must be reliable. In
addition to this, some of the use cases described in Section 3
require the interactions to be performed in a transactional fashion,
i.e. using a "request/response" mechanism. This is required so that
the network elements always remain in a known state, irrespective of
whether or not the transaction is successful.

2.4. **Interaction with Management Systems**

   When introducing an Access Node Control Mechanism, care is needed to
   ensure that the existing management mechanisms remain operational as
   before.

   Specifically when using the Access Node Control Mechanism for
   performing a configuration action on a network element, one gets
   confronted with the challenge of supporting multiple managers for the
   same network element: both the Element Manager as well as the Access
   Node Control Mechanism may now perform configuration actions on the
   same network element. Conflicts therefore need to be avoided.

   Using the Access Node Control Mechanism, the NAS retrieves and
   controls a number of Subscriber related parameters. The NAS may
   decide to communicate this information to a central Policy or AAA
   Server so that it kan keep track of the parameters and apply policies
   on them. The Server can then enforce those policies on the NAS. For
   instance, in case a subscriber is connected to more than one NAS, the
   Policy Server could be used to coordinate the bandwidth available on
   a given Access Port for use amongst the different NAS devices.

   Guidelines related to management will be addressed in Section 5.

2.5. **Circuit Addressing Scheme**

   In order to associate Subscriber parameters to a particular Access
   Port, the NAS needs to be able to uniquely identify the Access Port
   (or a specific circuit on an Access Port) using an addressing scheme.

   In deployments using an ATM aggregation network, the ATM PVC on an
   Access Loop connects the Subscriber to a NAS. Based on this
   property, the NAS typically includes a NAS-Port-Id, NAS-Port or
   Calling-Station-Id attribute in RADIUS authentication & accounting
   packets sent to the RADIUS server(s). Such attribute includes the
   identification of the ATM VC for this Subscriber, which allows in
   turn identifying the Access Loop.

   In an Ethernet-based aggregation network, a new addressing scheme is
   defined in [TR-101]. Two mechanisms can be used:

   o A first approach is to use a one-to-one VLAN assignment model for
      all Access Ports (e.g. a DSL port) and circuits on an Access Port
      (e.g. an ATM PVC on an ADSL port). This enables directly deriving
      the port and circuit identification from the VLAN tagging
      information, i.e. S-VLAN ID or <S-VLAN ID, C-VLAN ID> pair;

   o A second approach is to use a many-to-one VLAN assignment model
      and to encode the Access Port and circuit identification in the
      "Agent Circuit ID" sub-option to be added to a DHCP or PPPoE
      message. The details of this approach are specified in [TR-101].

   This document reuses the addressing scheme specified in TR-101. It
   should be noted however that the use of such a scheme does not imply
   the actual existence of a PPPoE or DHCP session, nor on the specific
   interworking function present in the Access Node. In some cases, no
   PPPoE or DHCP session may be present, while port and circuit
   addressing would still be desirable.

3. **Use Cases for Access Node Control Mechanism**

3.1. **Access Topology Discovery**

   [TR-059] and [TR-101] discuss various queuing/scheduling mechanisms
   to avoid congestion in the access network while dealing with multiple
   flows with distinct QoS requirements. One technique that can be used
   on a NAS is known as "Hierarchal Scheduling" (HS). This option is
   applicable in a single NAS scenario (in which case the NAS manages
   all the bandwidth available on the Access Loop) or in a dual NAS
   scenario (in which case the NAS manages some fraction of the Access
   Loop's bandwidth). The HS must, at a minimum, support 3 levels
   modelling the NAS port, Access Node uplink, and Access Loop sync
   rate. The rationale for the support of HS is as follows:

   o Provide fairness of network resources within a class.

   o Allow for a better utilization of network resources. Drop traffic
      early at the NAS rather than letting it traverse the aggregation
      network just to be dropped at the Access Node.

   o Enable more flexible Class of Service (CoS) behaviors other than
      only strict priority.

   o The HS system could be augmented to provide per application
      admission control.

   o Allow fully dynamic bandwidth partitioning between the various
      applications (as opposed to static bandwidth partitioning).

   o Support "per user weighted scheduling" to allow differentiated
      SLAs (e.g. business services) within a given traffic class.

   Such mechanisms require that the NAS gains knowledge about the
   topology of the access network, the various links being used and
   their respective rates. Some of the information required is somewhat
   dynamic in nature (e.g. DSL line rate - hence also the net data
   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 BRAS has of it.

   OSS systems are typically not designed to enforce the consistency of
   such data in a reliable and scalable manner across organizational
   boundaries. The Access Topology Discovery function is indended to
   allow the NAS to perform these functions without having to rely on an
   integration with an OSS system.

Communicating Access Loop attributes is specifically important in
case the rate of the Access Loop changes overtime. The DSL actual
data rate may be different every time the DSL NT is turned on. In
this case, the Access Node sends an Information Report message to the
NAS after the DSL line has resynchronized.

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

The hierarchy and the rates of the various links to enable the NAS
hierarchical scheduling and policing mechanisms are the following:

o The identification and speed (data rate) of the DSL Access Loop
   (i.e. the net data rate)

o The identification and speed (data rate) of the Remote
   Terminal(RT)/Access Node uplink (when relevant)

The NAS can adjust downstream shaping to current Access Loop actual
data rate, and more generally re-configure the appropriate nodes of
its hierarchical scheduler (support of advanced capabilities
according to TR-101).

This use case may actually include more information than link
identification and corresponding data rates. In case of DSL Access
Loops, the following Access Loop characteristics can be sent to the
NAS (cf. ITU-T Recommendation G.997.1 [G.997.1]):

o DSL Type (e.g. ADSL1, ADSL2, SDSL, ADSL2+, VDSL, VDSL2)

o Framing mode (e.g. ATM, ITU-T Packet Transfer Mode (PTM), IEEE
   802.3 Ethernet in the First Mile (EFM))

o DSL port state (e.g. synchronized/showtime, low power, no power/
   idle)

o Actual net data rate (upstream/downstream)

o Maximum achievable/attainable net data rate (upstream/downstream)

o Minimum net data rate configured for the Access Loop (upstream/
   downstream)

o Maximum net data rate configured for the Access Loop (upstream/
   downstream)

o Minimum net data rate in low power state configured for the Access
   Loop (upstream/downstream)

o Maximum achievable interleaving delay (upstream/downstream)

o Actual interleaving delay (upstream/downstream)

The NAS MUST be able to receive Access Loop characteristics
information, and share such information with AAA/Policy Servers.

3.2. Access Loop Configuration

Access Loop rates are typically configured in a static way. When a
Subscriber wants to change its Access Loop rate, the network operator
needs to reconfigure the Access Port configuration, possibly implying
a business-to-business transaction between an Internet Service
Provider (ISP) and an Access Provider. From an Operating
Expenditures (OPEX) perspective this is a costly operation.

Using the Access Node Control Mechanism to change the Access Loop
rate from the NAS avoids those cross-organization business-to-
business interactions and allows to centralize Subscriber-related
service data in e.g. a Policy Server. More generally, several Access
Loop parameters (e.g. minimum data rate, interleaving delay) could be
changed by means of the Access Node Control Mechanism.

Triggered by the communication of the Access Loop attributes
described in Section 3.1, the NAS could query a Policy or AAA Server
to retrieve Access Loop configuration data. The best way to change
Access Loop parameters is by using profiles. These profiles (e.g.
DSL profiles for different services) are pre-configured by the
Element Manager managing the Access Nodes. The NAS may then use the
Configure Request message to send a reference to the right profile to
the Access Node. The NAS may also update the Access Loop
configuration due to a Subscriber service change (e.g. triggered by
the Policy Server).

The Access Loop Configuration mechanism may also be useful for
configuration of parameters that are not specific to the Access Loop
technology. Examples include the QoS profile to be used for an
Access Loop, or the per-Subscriber multicast channel entitlement
information, used for IPTV applications where the Access Node is
performing IGMP snooping or IGMP proxy function. The latter is also
discussed in Section 3.4.

It may be possible that a Subscriber wants to change its Access Loop
rate, and that the operator wants to enforce this updated Access Loop
Rate on the Access Node using ANCP, but that the Access Node Control
Adjacency is down. In such a case, the NAS will not be able to
request the configuration change on the Access Node. The NAS should
then report this failure to the external management system, which
could use application specific signaling to notify the Subscriber of
the fact that the change could not be performed at this time.

3.3. Remote Connectivity Test

Traditionally, ATM circuits are point to point connections between
the BRAS and the DSLAM or DSL NT. In order to test the connectivity
on layer 2, appropriate OAM functionality is used for operation and
troubleshooting. An end-to-end OAM loopback is performed between the
edge devices (NAS and HGW) of the broadband access network.

When migrating to an Ethernet-based aggregation network (as defined
by TR-101), end to end ATM OAM functionality is no longer applicable.
Ideally in an Ethernet aggregation network, end-to-end Ethernet OAM
as specified in IEEE 802.1ag and ITU-T Recommendation Y.1730/1731 can
provide Access Loop connectivity testing and fault isolation.
However, most HGWs do not yet support these standard Ethernet OAM
procedures. Also, various access technologies exist such as ATM/DSL,
Ethernet in the First Mile (EFM) etc. Each of these access
technologies have their own link-based OAM mechanisms that have been
or are being standardized in different standard bodies.

In a mixed Ethernet and ATM access network (including the local
loop), it is desirable to keep the same ways to test and troubleshoot
connectivity as those used in an ATM based architecture. To reach
consistency with the ATM based approach, an Access Node Control
Mechanism between NAS and Access Node can be used until end-to-end
Ethernet OAM mechanisms are more widely available.

Triggered by a local management interface, the NAS can use the Access
Node Control Mechanism to initiate an Access Loop test between Access
Node and HGW. In case of an ATM based Access Loop the Access Node
Control Mechanism can trigger the Access Node to generate ATM (F4/F5)
loopback cells on the Access Loop. In case of Ethernet, the Access
Node can perform a port synchronization and administrative test for
the Access Loop. The Access Node can send the result of the test to
the NAS via a Control Response message. The NAS may then send the
result via a local management interface. Thus, the connectivity
between the NAS and the HGW can be monitored by a single trigger
event.

3.4. Multicast

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

   In case of an ATM access/aggregation network, such as the reference
   architecture specified in Broadband Forum [TR-059], multicast traffic
   replication is performed in the NAS. In this model, typically IGMP
   is used to control the multicast replication process towards the
   subscribers. The NAS terminates and processes IGMP signaling
   messages sent by the subscribers; towards the Regional Network, the
   NAS typically uses a multicast routing protocol such as PIM. The ATM
   Access Nodes and aggregation switches don't perform IGMP processing,
   nor do they perform multicast traffic replication. As a result,
   network resources are wasted within the access/aggregation network.

   To overcome this resource inefficiency, 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.

   Another option is for the subscriber to communicate the "join/leave"
   information with the NAS. This can for instance be done by
   terminating all subscriber IGMP 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 can use ANCP to create replication state in the AN
   for efficient multicast replication. The NAS sends a single copy of
   the multicast stream towards the AN. The NAS can perform conditional
   access and multicast admission control on multicast joins, and create
   replication state in the AN if the flow is admitted by the NAS.

   The following subsections describe the different use cases related to

multicast.

3.4.1. Multicast Conditional Access

In a DSL Broadband access scenario Service Providers may want to
dynamically control, at the network level, access to some multicast
flows on a per user basis. This may be used in order to
differentiate among multiple Service Offers or to realize/reinforce
conditional access for sensitive content. Note that, in some
environments, application layer conditional access by means of
Digital Rights Management (DRM) may provide sufficient control, so
that Multicast Conditional Access may not be needed.

Where Multicast Conditional Access is required, it is possible, in
some cases, to provision the necessary conditional access information
into the AN so the AN can then perform the conditional access
decisions autonomously. For these cases, the NAS can use ANCP to
provision the necessary information in the AN so that the AN can
decide locally to honor a join or to not honor a join. This can be
done with the Control Request and Control Response messages.

Provisioning the conditional access information on the AN can be done
using a "White list", "Grey list", and/or a "Black list". A White
list associated with an Access Port identifies the multicast flows
that are allowed to be replicated to that port. A Black list
associated with an Access Port identifies the multicast flows that
are not allowed to be replicated to that port. A Grey list
associated with an Access Port identifies the multicast flows for
which the AN on receiving a join message, before starting traffic
replication queries the NAS for further authorization. Each list
contains zero, one or multiple entries and each entry may specify a
single flow or contain ranges (i.e. mask on Group address and/or mask
on Source address).

Upon receiving a join message on an Access Port, the Access Node will
first check if the requested multicast flow is part of a White, Grey
or a Black list associated with that Access Port. If it is part of a
White list, the AN autonomously starts replicating multicast traffic.
If it is part of a Black list, the AN autonomously discards the
message because the request is not authorized, and may thus inform
the NAS and log the request accordingly. If it is part of a Grey
list the AN uses ANCP to query the NAS, that in turn will respond to
the AN indicating whether the join is to be honored (and hence
replication performed by the AN) or denied (and hence replication not
performed by the AN.)

If the requested multicast flow is part of multiple lists associated
with the Access Port, then the most specific match will be used. If

the most specific match occurs in multiple lists, the Black list
entry takes precedence over the Grey list, which takes precedence
over the White list.

If the requested multicast flow is not part of any list, the message
should be discarded. This default behavior can easily be changed by
means of a "catch-all" statement in either the White list or the Grey
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. Similarly, if the default
behavior should be to send a request to the NAS, then adding
(<S=*,G=*>) in the Grey List accomplishes that.

The White List, Black List and Grey List can contain entries
allowing:

o   an exact match for a (*,G) ASM group (e.g. <G=g.h.i.l>);

o   an exact match for a (S,G) SSM channel (e.g.
    <S=s.t.u.v,G=g.h.i.l>);

o   a mask-based range match for a (*,G) ASM group (e.g. <G=g.h.i.l/
    Mask>);

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

Following are some example configurations:

o   Scenario 1: reject all messages

    *   Black List = {<S=*,G=*>}

o   Scenario 2: reject all messages, except Join (S=*,G=Gi) (1<=i<=n)

    *   White List = { <S=*,G=G1> , <S=*,G=G2>, ... <S=*,G=Gn>}

    *   Black List = {<S=*,G=*>}

o   Scenario 3: AN performs autonomous decisions for some channels,
    and asks the NAS for other channels

    *   White List = { <S=*,G=G1> , <S=*,G=G2>, ... <S=*,G=Gn>}

    *   Grey List = { <S=s,G=Gm>} for m>n

    *   Black list = {<S=*,G=*>}

    * ==> Join (S=*,G=Gi) gets honored by AN (1<=i<=n)

    * ==> Join (S=s,G=Gm) triggers ANCP Admission Request to NAS

    * ==> everything else gets rejected by AN

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

In some other cases, it may be desirable to have the conditional access decision being taken by the NAS or a Policy Server. This may be the case when conditional access information changes frequently, or when the multicast groups are not known to a client application in advance. The conditional access control could be tied to a more complex policy/authorization mechanism, e.g. time of day access, or location based access or to invoke a remote authorisation server. For these cases, the AN can use ANCP to query the NAS, that in turn will respond to the AN indicating whether the join is to be honored (and hence replication performed by the AN) or denied. This can be done with the Admission Request and Admission Response messages.

Some examples of using NAS querying are the following:

o Roaming users: a subscriber that logs in on different wireless
   hotspots and would like to receive multicast content he is
   entitled to receive;

o A related example is that of mobility or seamless handover; in
   both cases the burden of (re)configuring access nodes with White
   lists or Black lists may be too high;

o "Over The Top Video Partnerships": Service Providers may choose to
   partner with Internet video providers to provide video content.
   In this case, the multicast group mappings may not be known in
   advance, or may be reused for different content in succession.

o "Pay Per View": This is the case when a subscriber chooses a
   specific IPTV channel which is made available for a given amount
   of time.

3.4.2. Multicast Admission Control

The successful delivery of Triple Play Broadband services is quickly
becoming a big capacity planning challenge for most of the Service
Providers nowadays. Solely increasing available bandwidth is not
always practical, cost-economical and/or sufficient to satisfy end
user experience given not only the strict requirements of unicast
delay sensitive applications like VoIP and Video, but also the fast
growth of multicast interactive applications such as
videoconferencing, digital TV, digital audio, online movies and
networked gaming. These applications are typically characterized by
a delay sensitive nature, an extremely loss sensitive nature and
intensive bandwidth requirements. They are also typically "non-
elastic", which means that they operate at a fixed bandwidth, that
cannot be dynamically adjusted to the currently available bandwidth.

Therefore a Connection Admission Control (CAC) mechanism covering
admission of video traffic over the DSL Broadband access is required,
in order to avoid oversubscribing the available bandwidth and
negatively impacting the end user experience.

Considering specifically admission control over the access line,
before honoring a user request to join a new multicast flow, the
combination of AN and NAS must ensure admission control is performed
to validate that there is sufficient bandwidth remaining on the
access line to carry the new video stream (in addition to all other
multicast and unicast video streams sent over the access line). The
solution needs to cope with multiple flows per access line and needs
to allow access line bandwidth to be dynamically shared across
multicast and unicast traffic (the unicast CAC is performed either by
the NAS or by some off-path Policy Server).

Thus, supporting CAC for the access line requires some form of
synchronization between the entity performing multicast CAC (e.g. the
NAS or the AN), the entity performing unicast CAC (e.g. the Policy
Server) and the entity actually enforcing the multicast replication
(i.e. the AN). This synchronization can be achieved in a number of
ways:

o One approach is for the AN to query the NAS so that Admission
   Control for the access line is performed by the NAS, or by the
   Policy Server which interacts with the AN via NAS. The AN can use
   ANCP to query the NAS, that in turn performs multicast Admission
   Control check for the new multicast flow and responds to the AN
   indicating whether the join is to be honored (and hence
   replication performed by the AN) or denied. The NAS may locally
   keep track of the portion of the Access Loop net data rate that is
   available for (unicast or multicast) video flows, and perform

video bandwidth accounting for the Access Loop. Upon receiving an
Admission Request from the AN, the NAS can check available Access
Loop bandwidth before admitting or denying the multicast flow. In
the process, the NAS may communicate with the Policy Server. For
unicast video services such as Video on Demand (VoD), the NAS may
also be queried (by a Policy Server or via on-path CAC signaling),
so that it can perform admission control for the unicast flow, and
update the remaining available Access Loop bandwidth. The ANCP
requirements to support this approach are specified in this
document.

o The above model could be enhanced with the notion of "Delegation
   of Authorization". In such a model, the NAS or the Policy Server
   delegates authority to the Access Node to perform multicast
   Admission Control on the Access Loop. This is sometimes referred
   to as "Bandwidth Delegation", referring to the portion of the
   total Access Loop bandwidth which can be used by the Access Node
   for multicast Admission Control. In this model, the NAS or the
   Policy Server manages the total access line bandwidth, performs
   unicast admission control and uses ANCP to authorize the Access
   Node to perform multicast Admission Control within the bounds of
   the "delegated bandwidth". Upon receiving a request for a
   multicast flow replication which matches an entry in the White or
   Grey List, the AN performs the necessary bandwidth admission
   control check for the new multicast flow, before starting the
   multicast flow replication. At this point, there is typically no
   need for the Access Node to communicate with the NAS or the Policy
   Server via the NAS. The ANCP requirements to support this
   approach are also specified in this document.

o In case the subscriber communicates the "join/leave" information
   with the NAS (e.g. by terminating all subscriber IGMP signaling on
   the NAS, or by using some form of application level signaling),
   the approach is very similar. In this case, the NAS may locally
   keep track of the portion of the Access Loop bandwidth that is
   available for video flows, perform CAC for unicast and multicast
   flows and perform video bandwidth management. The NAS can set the
   replication state on the AN using ANCP if the flow is admitted.
   For unicast video services the NAS may be queried (by a policy
   server or via on-path CAC signaling) to perform admission control
   for the unicast flow, and update the remaining available Access
   Loop bandwidth. The ANCP requirements to support this approach
   are specified in this document;

o In the last approach, the Policy Server queries the AN directly or
   indirectly via the NAS, so that both unicast and multicast CAC for
   the access line are performed by the AN. In this case, a
   subscriber request for a unicast flow (e.g. a Video on Demand

      session) will trigger a resource request message towards a Policy
      Server; the latter will then query the AN (possibly via the NAS),
      that in turn will perform unicast CAC for the access line and
      respond, indicating whether the unicast request is to be honored
      or denied. The above model could also be enhanced with the notion
      of "Delegation of Authorization". In such a model, the Policy
      Server delegates authority to the Access Node to perform multicast
      Admission Control on the Access Loop. In the case when the Policy
      Server queries the AN directly, the approach doesn't require the
      use of ANCP. It is therefore beyond the scope of this document.
      In the case when the Policy Server queries the AN indirectly via
      the NAS, the approach requires the use of ANCP and is therefore in
      the scope of this document.

3.4.2.1. **Delegation of Authority - Bandwidth Delegation**

   The NAS uses ANCP to indicate to the AN whether or not Admission
   Control is required for a particular multicast flow on a given Access
   Port. In case Admission Control is required, the Access Node needs
   to know whether or not it is authorized to perform Admission Control
   itself and, if so, within which bounds it is authorized to do so
   (i.e. how much bandwidth is "delegated" by the NAS or the Policy
   Server). Depending on the type of multicast flow, Admission Control
   may or may not by done by the AN:

   o Multicast flows that require a Conditional Access operation to be
      performed by the Access Node, are put in the Black or White List.
      In addition, the Access Node performs Admission Control for those
      flows in the White List for which it is authorized to do so.

   o Multicast flows that require a Conditional Access operation to be
      performed by the NAS or the Policy Server, are put in the Grey
      List. In addition, for those flows in the Grey List for which the
      Access Node should perform Admission Control, the NAS or the
      Policy Server will delegate authority to the AN.

   In some cases, the bandwidth that the NAS or the Policy Server
   initially delegated to the AN may not be enough to satisfy a
   multicast request for a new flow. In this scenario the AN can use
   ANCP to query the NAS in order to request additional delegated
   multicast bandwidth. This is a form of extending the AN
   authorization to perform Admission Control. The NAS or the Policy
   Server decides if the request for more bandwidth can be satisfied and
   uses ANCP to send a response to the AN indicating the updated
   delegated multicast bandwidth. It is worth noting that in this case,
   the time taken to complete the procedure is an increment to the
   zapping delay: in order to minimize the zapping delay for future join
   requests the AN can insert in the request message two values: the

   minimum amount of additional multicast bandwidth requested, and the
   preferred additional amount. The first value is the amount that
   allows the present join request to be satisfied, the second value an
   amount that anticipates further join requests.

   In some cases, the NAS or the Policy Server may not have enough
   unicast bandwidth to satisfy a new incoming video request: in these
   scenarios the NAS can use ANCP to query (or instruct) the AN in order
   to decrease the amount of multicast bandwidth previously delegated on
   a given Access Port. This is a form of limiting/withdrawing AN
   authorization to perform Admission Control. The NAS can use ANCP to
   send a response to AN indicating the updated delegated multicast
   bandwidth. Based on considerations similar to those of the previous
   paragraph, it indicates the minimum amount of multicast bandwidth
   that it needs released and a preferred amount, which may be larger.

   Note: in order to avoid impacting existing multicast traffic, the NAS
   must not decrease the amount of delegated multicast bandwidth to a
   value lower than the bandwidth that is currently in use. This
   requires the NAS to be aware of this information (e.g. by means of a
   separate query action).

   In addition, in some cases, upon receiving a leave for a specific
   multicast flow, the AN may decide that it has an excess of delegated
   but uncommitted bandwidth. In such case, the AN can use ANCP to send
   a message to the NAS to release (part of the) unused multicast
   bandwidth previously delegated. In this process, the Access Node may
   decide to retain a minimum amount of bandwidth for multicast
   services.

3.4.2.2. When not to perform Admission Control for a subset of flows

   In general, the Access Node and NAS may not be aware of all possible
   multicast groups that will be streamed in the access network. For
   instance, it is likely that there will be multicast streams offered
   across the Internet. For these unknown streams, performing bandwidth
   Admission Control may be challenging.

   To solve this, these requests could be accepted without performing
   Admission Control. This solution works, provided that the network
   handles the streams as best effort, so that other streams (that are
   subject to Admission Control) are not impacted at times of
   congestion.

   Disabling Admission Control for an unknown stream can be achieved by
   adding a "catch-all statement" in the Access Node white list or grey
   list. In case the Access Node queries the NAS, the NAS on his turn
   will have to accept the request. That way, the unknown streams are

not blocked by default.

Next, in order to ensure that the streams are handled as best effort, the flow must be marked as such when entering the service provider network. This way, whenever congestion occurs somewhere in the access/aggregation network, this stream will be kicked out before the access provider's own premium content.

The above concept is applicable beyond the notion of "Internet streams" or other unknown streams; it can applied to known multicast streams as well. In this case, the Access Node or NAS will accept the stream even when bandwidth may not be sufficient to support the stream. This again requires that the stream is marked as best effort traffic before entering the access/aggregation network.

3.4.2.3. Multicast Admission Control and White Lists

As mentioned in section Section 3.4.1, conditional access to popular IPTV channels can be achieved by means of a White and Black list configured on the Access Node. This method allows the Access Node to autonomously decide whether or not access can be granted to a multicast flow.

IPTV is an example of a service that will not be offered as best effort, but requires some level of guaranteed Quality of Service. This requires the use of Multicast Admission Control. Hence, if the Access Node wants to autonomously perform the admission process, it must be aware of the bandwidth characteristics of multicast flows. Otherwise, the Access Node would have to query the NAS for Multicast Admission Control (as per the Grey list behavior); this would defeat the purpose of using a White and Black list.

Some network deployments may combine the use of White list, Black list and Grey list. The implications of such a model to the overall Multicast Admission Control model are not fully explored in this document.

3.4.3. Multicast Accounting and Reporting

It may be desirable to perform time and/or volume based accounting for certain multicast flows sent on particular Access Ports. In case the AN is performing the traffic replication process, it knows when replication of a multicast flow to a particular Access Port or user start and stops. Multicast accounting can be addressed in two ways:

o The AN keeps track of when replication for a given multicast flow starts or ends on a specified Access Port, and generates time and/or volume based accounting information per Access Port and per

multicast flow, before sending it to a central accounting system
for logging. Given that the AN communicates with the accounting
system directly, the approach doesn't require the use of ANCP. It
is therefore beyond the scope of this document;

o The AN keeps track of when replication for a given multicast flow
   starts or ends on a specified Access Port, and reports this
   information to the NAS for further processing. In this case, ANCP
   can be used to send the information from the AN to the NAS. This
   will be discussed in the remainder of this document.

The Access Node can send multicast accounting information to the NAS
using the Information Report message. A distinction can be made
between two cases:

o Basic accounting information: the Access Node informs the NAS
   whenever replication starts or ends for a given multicast flow on
   a particular Access Port;

o Detailed accounting information: the Access Node not only informs
   the NAS when replication starts or ends, but also informs the NAS
   about the multicast traffic volume replicated on the Access Port
   for that multicast flow. This is done by adding a byte count in
   the Information Report message which is sent to the NAS when
   replication ends.

Upon receiving the Information Report messages, the NAS generates the
appropriate time and/or volume based accounting records per Access
Loop and per multicast flow, to be sent to the accounting system.

The NAS should inform the Access Node about the type of accounting
needed for a given multicast flow on a particular Access Port:

o No reporting messages need to be sent to the NAS

o Basic accounting is required

o Detailed accounting is required

Note that in case of very fast channel changes, the amount of
Information Report messages to be sent to the NAS could become high.

The ANCP requirements to support this use case are specified below in
this document.

It may also be desirable for the NAS to have the capability to
asynchronously query the AN to obtain an instantaneous status report
related to multicast flows currently replicated by the AN. Such a

reporting functionality could be useful for troubleshooting and
monitoring purposes. The NAS can query the AN to know the following:

o Which flows are currently being sent on a specific Access Port
   (i.e. a report for one Access Port)

o On which Access Ports is a specified multicast flow currently
   being sent (i.e. a report for one multicast flow)

o Which multicast flows are currently being sent on each of the
   Access Ports (i.e. a global report for one Access Node)

3.4.4. Spontaneous Admission Response

The capability to dynamically stop the replication of a multicast
flow can be useful in different scenarios: for example in case of
prepaid service, when available credit expires, the Service Provider
may want to be able to stop multicast replication on a specified
Access Port for a particular user. Another example of applicability
for this functionality is a scenario where a Service Provider would
like to show a "Content Preview": in this case a multicast content
will be delivered just for a fixed amount of time.

In both cases an external entity (for example a Policy Server or an
external application entity), can instruct the NAS to interrupt the
multicast replication of a specified multicast flow to a specified
Access Port or user. The NAS can then use ANCP to communicate this
decision to the Access Node. This can be done with the Admission
Response message.

In some deployment scenarios, the NAS may be made aware of end-users
requests to join/leave a multicast flow by other means than ANCP
Admission Requests sent by the AN. One possible deployment scenario
where this model applies is the case where the Access Node doesn't
process the IGMP join/leave messages from the end-user (e.g. because
they are tunneled), but forwards them to the NAS. In such
environments, the NAS can control multicast replication on the AN via
ANCP through the use of Spontaneous Admission Responses (i.e. sent by
the NAS without prior receipt of a corresponding Admission Request).

4. Requirements

4.1. ANCP Functional Requirements

   R-1   The ANCP MUST be easily extensible through the definition of new
         message types or TLVs to support use cases beyond those
         currently addressed in this document (this includes the use of
         Access Nodes different from a DSLAM, e.g. a PON Access Node).

   R-2   The ANCP MUST be flexible enough to accommodate the various
         technologies that can be used in an access network and in the
         Access Node; this includes both ATM and Ethernet.

   R-3   The Access Node Control interactions MUST be reliable (using
         either a reliable transport protocol (e.g. TCP) for the Access
         Node Control Protocol messages, or by designing ANCP to be
         reliable).

   R-4   The ANCP MUST support "request/response" transaction-based
         interactions for the NAS to communicate control decisions to the
         Access Node, or for the NAS to request information from the
         Access Node. Transactions MUST be atomic, i.e. they are either
         fully completed, or rolled-back to the previous state. This is
         required so that the network elements always remain in a known
         state, irrespective of whether or not the transaction is
         successful.

   In case the NAS wants to communicate a bulk of independent control
   decisions to the Access Node, the transaction (and notion of
   atomicity) applies to the individual control decisions. This avoids
   having to roll back all control decisions. Similarly, if the NAS
   wants to request a bulk of independent information elements from the
   Access Node, the notion of transaction applies to the individual
   information elements.

   R-5   The ANCP MUST be scalable enough to allow a given NAS to control
         at least 5000 Access Nodes.

   R-6   The operation of the ANCP in the NAS and Access Nodes MUST be
         controllable via a management station (e.g. via SNMP). This
         MUST allow a management station to retrieve statistics and
         alarms related to the operation of the ANCP, as well as to allow
         it to initiate OAM operations and retrieve corresponding
         results.

4.2. ANCP Multicast Requirements

R-7   The ANCP MUST support providing multicast conditional access
       information to Access Ports on an Access Node, using Black,
       Grey and White lists.

R-8   The ANCP MUST support binding a particular Black, Grey and
       White List to a given Access Port.

R-9   Upon receiving a join to a multicast flow which matches the
       Grey List the ANCP MUST allow the AN to query the NAS to
       request an admission decision for replicating that multicast
       flow to a particular Access Port.

R-10   The ANCP MUST allow the NAS to send an admission decision to
       the AN indicating whether or not a multicast flow may be
       replicated to a particular Access Port.

R-11   The ANCP MUST allow the NAS to indicate to the AN whether or
       not Admission Control is needed for some multicast flows on a
       given Access Port and where needed whether or not the Access
       Node is authorized to perform Admission Control itself (i.e.
       whether or not AN Bandwidth Delegation applies).

R-12   In case of Admission Control without AN Bandwidth Delegation,
       the ANCP MUST allow the NAS to reply to a query from the AN
       indicating whether or not a multicast flow is allowed to be
       replicated to a particular Access Port.

R-13   In case of Admission Control with AN Bandwidth Delegation, the
       ANCP MUST allow the NAS to delegate a certain amount of
       bandwidth to the AN for a given Access Port for multicast
       services only.

R-14   In case of Admission Control with AN Bandwidth Delegation, the
       ANCP MUST allow the AN to query the NAS to request additional
       multicast bandwidth on a given Access Port.

R-15   In case of Admission Control with AN Bandwidth Delegation, the
       ANCP MUST allow the NAS to query (or to instruct) the AN to
       reduce the amount of bandwidth previously delegated on a given
       Access Port.

R-16   In case of Admission Control with AN Bandwidth Delegation, the
       ANCP MUST allow the AN to inform the NAS in case it
       autonomously releases redundant multicast bandwidth on a given
       Access Port.

   R-17   The ANCP MUST allow the AN to send an Information Report
          message to the NAS whenever replication of a multicast flow on
          a particular Access Port start or ends.

   R-18   The ANCP MUST allow the AN to send an Information Report
          message to the NAS indicating the multicast traffic volume that
          has been replicated on that port.

   R-19   The ANCP MUST allow the NAS to indicate to the AN whether or
          not multicast accounting is needed for a multicast flow on a
          particular Access Port.

   R-20   In case multicast accounting is needed for a multicast flow on
          a particular Access Port, the ANCP MUST allow the NAS to
          indicate to the AN whether or not additional volume accounting
          information is required.

   R-21   The ANCP MUST allow the NAS to revoke a decision to replicate a
          multicast flow to a particular Access Port, which had been
          conveyed earlier to an AN.

   R-22   The ANCP MUST support partial updates of the White, Grey and
          Black lists.

   R-23   The ANCP MUST allow the NAS to query the AN to obtain
          information on what multicast flows are currently being
          replicated on a given Access Port, what Access Ports are
          currently receiving a given multicast flow, or what multicast
          flows are currently replicated on each Access Port.

4.3. Protocol Design Requirements

   R-24   The ANCP SHOULD provide a "shutdown" sequence allowing the
          protocol to inform the peer that the system is gracefully
          shutting down.

   R-25   The ANCP SHOULD include a "report" model for the Access Node to
          spontaneously communicate to the NAS changes of states.

   R-26   The ANCP SHOULD support a graceful restart mechanism to enable
          it to be resilient to network failures between the AN and NAS.

   R-27   The ANCP MUST provide a means for the AN and the NAS to inform
          each peer about the supported use cases (either use cases
          defined in this document or future use cases yet to be
          defined), and to negotiate a common subset.

4.4. Access Node Control Adjacency Requirements

The notion of an Access Node Control Adjacency is defined in
Section 1.2.

R-28   The ANCP MUST support an adjacency protocol in order to
       automatically synchronize its operational state between its
       peers, to agree on which version of the protocol to use, to
       discover the identity of its peers, and detect when they
       change.

R-29   The ANCP MUST include a mechanism to automatically detect
       adjacency loss.

R-30   A loss of the Access Node Control Adjacency MUST NOT affect
       subscriber connectivity.

R-31   If the Access Node Control Adjacency is lost, it MUST leave the
       network elements in a known state, irrespective of whether or
       not the ongoing transaction was successful.

R-32   The ANCP MUST support a mechanism to synchronize access port
       configuration and status information between ANCP peers as part
       of establishing or recovering the Access Node Control
       Adjacency.

4.5. ANCP Transport Requirements

R-33   The Access Node Control Mechanism MUST be defined in a way that
       is independent of the underlying layer 2 transport technology.
       Specifically, the Access Node Control Mechanism MUST support
       transmission over an ATM as well as over an Ethernet
       aggregation network.

R-34   The ANCP MUST use the IP protocol stack.

R-35   If the layer 2 transport technology is based on ATM, then the
       ANCP peers must use the encapsulation according to [RFC2684]
       routed (IPoA).

R-36   If the layer 2 transport technology is based on Ethernet, then
       the ANCP peers must use the encapsulation according to [RFC894]
       (IPoE).

4.6. Access Node Requirements

   This section lists the requirements for an AN that supports the use
   cases defined in this document. Note that this document does not
   intend to impose absolute requirements on network elements.
   Therefore, the words "must" and "should" used in this section are not
   capitalized.

4.6.1. General Architecture

   The Access Node Control Mechanism is defined to operate between an
   Access Node (AN) and a NAS. In some cases, one AN can be connected
   to more than one physical NAS devices (e.g. in case different
   wholesale service providers have different NAS devices). In such a
   model, the physical AN needs to be split in virtual ANs, each having
   its own Access Node Control reporting and/or enforcement function.

   R-37 An Access Node as physical device can be split in logical
         partitions. Each partition may have its independent NAS.
         Therefore the Access Node must support at least 2 partitions.
         The Access Node should support 8 partitions.

   R-38 One partition is grouped of several Access Ports. Each Access
         Port on an Access Node must be assigned uniquely to one
         partition.

   It is assumed that all circuits (i.e. ATM PVCs or Ethernet VLANs) on
   top of the same physical Access Port are associated with the same
   partition. In other words, partitioning is performed at the level of
   the physical Access Port only.

   R-39 Each AN partition must have a separate Access Node Control
         Adjacency to a NAS.

   R-40 Each AN partition must be able to enforce access of the
         controllers to their designated partitions.

   R-41 The Access Node should be able to establish and maintain ANCP
         Adjacencies to redundant controllers.

4.6.2. Control Channel Attributes

   The Control Channel is a bidirectional IP communication interface
   between the controller function (in the NAS) and the reporting/
   enforcement function (in the AN). It is assumed that this interface
   is configured (rather than discovered) on the AN and the NAS.

   Depending on the network topology, the Access Node can be located in

a street cabinet or in a central office. If an Access Node in a
street cabinet is connected to a NAS, all user traffic and Access
Node Control data can use the same physical link.

R-42 The Control Channel should use the same facilities as the ones
      used for the data traffic. Note that this is actually a
      deployment consideration, which has no impact on the actual
      protocol design.

R-43 The Access Node must process control transactions in real-time
      (i.e. with a specific response latency).

R-44 The Access Node should mark Access Node Control Protocol
      messages with a high priority (e.g. VBR-rt for ATM cells,
      p-bit 6 or 7 for Ethernet packets) in order to avoid or reduce
      the likelihood of dropping packets in case of network
      congestion.

R-45 If ATM interfaces are used then any VPI and VCI value must be
      able to be used for the purpose of supporting the Access Node
      Control Channel.

R-46 If Ethernet interfaces are used then any C-VID and S-VID must
      be able to be used for the purpose of supporting the Access
      Node Control Channel.

4.6.3. Capability Negotiation Failure

R-47 In case the Access Node and NAS cannot agree on a common set of
      capabilities, as part of the ANCP capability negotiation
      procedure, the Access Node must report this to network
      management.

4.6.4. Adjacency Status Reporting

R-48 The Access Node should support generating an alarm to a
      management station upon loss or malfunctioning of the Access
      Node Control Adjacency with the NAS.

4.6.5. Identification

R-49 To identify the Access Node and Access Port within a control
      domain a unique identifier is required. This identifier must
      be in line with the addressing scheme principles specified in
      section 3.9.3 of TR-101.

   R-50 In a Broadband Forum TR-101 network arcitecture, an Access
         Circuit Identifier (ACI) identifying an AN and Access Port is
         added to DHCP and PPPoE messages. The NAS must use the same
         ACI format in ANCP messages in order to allow the NAS to
         correlate this information with the information present in DHCP
         and PPPoE messages.

4.6.6. Multicast

   R-51 The AN must deny any join to a multicast flow matching the
         Black List for the relevant Access Port.

   R-52 The AN must accept any join to a multicast flow matching the
         White List and for which no Bandwidth Delegation is used.

   R-53 Upon receiving a join to a multicast flow which matches the
         White List and for which Bandwidth Delegation is used, the AN
         must perform the necessary bandwidth admission control check
         for the new flow before starting the multicast flow
         replication. This may involve a decision taken locally, or
         querying the NAS or external system such as a Policy Server, to
         request additional delegated multicast bandwidth on a given
         Access Port.

   R-54 Upon receiving a join to a multicast flow which matches the
         Grey List and for which no Bandwidth Delegation is used, the AN
         must support using ANCP to query the NAS to receive a response
         indicating whether that join is to be honored or denied. In
         this case, the NAS will perform both the necessary conditional
         access and the admission control checks for the new flow.

   R-55 Upon receiving a join to a multicast flow which matches the
         Grey List and for which Bandwidth Delegation is used, the AN
         must first perform the necessary bandwidth admission control
         check for the new flow. If successful, the AN must support
         using ANCP to query the NAS to receive a response indicating
         whether that join is to be honored or denied.

   R-56 In case of Admission Control with AN Bandwidth Delegation, the
         AN must support using ANCP to notify the NAS when the the user
         leaves the multicast flow.

   R-57 In case of Admission Control with AN Bandwidth Delegation, the
         AN must support using ANCP to query the NAS to request
         additional delegated multicast bandwidth on a given Access
         Port; the AN should be able to specify both the minimum and the
         preferred amount of additional multicast bandwidth requested.

R-58   In case of Admission Control with AN Bandwidth Delegation, upon
        receiving a Bandwidth Delegation Request from the NAS querying
        the AN for the delegated multicast bandwidth on a given Access
        Port, the AN must support using ANCP to send a Bandwidth
        Delegation Response, indicating the currently delegated
        multicast bandwidth.

R-59   In case of Admission Control with AN Bandwidth Delegation, it
        may happen that the NAS wants to "revoke" (part of) the
        delegated bandwidth. Part of the previously delegated
        bandwidth may however be in use by multicast services.
        Therefore, upon receiving a Bandwidth Delegation Request from
        the NAS instructing to decrease the delegated multicast
        bandwidth on a given Access Port, the AN must support using
        ANCP to send a Bandwidth Delegation Response, indicating the
        delegated multicast bandwidth after the decrease (indicating
        how much of the delegated bandwidth can be returned to the NAS
        without impacting multicast services that are currently
        running).

R-60   In case of Admission Control with AN Bandwidth Delegation, the
        AN must support using ANCP to send a Bandwidth Release message
        to the NAS in order to release unused delegated multicast
        bandwidth on a given Access Port.

R-61   If the requested multicast flow is not part of any list
        associated with the Access Port, the AN must discard the
        message.

R-62   If the requested multicast flow is part of multiple lists
        associated with the Access Port, the AN must use the most
        specific match.

R-63   If the requested multicast flow has the same most specific
        match in multiple lists, the AN must give precedence to the
        Black list, followed by the Grey list, and then the White list.

R-64   The AN must support configuring a "catch-all" statement in the
        Black, White or Grey list in order to enforce a default
        behavior for a join to a multicast flow which doesn't match any
        other entry in a list for the relevant Access Port.

R-65   Upon querying the NAS, the AN must not propagate the join
        message before the successful authorization from the NAS is
        received.

R-66   Upon receiving a leave for a multicast flow which matches the
       Grey List, the AN should be able to autonomously stop
       replication and advertise this event to the NAS.

R-67   The AN must support using ANCP to send an Information Report
       message to the NAS whenever replication starts or ends.

R-68   The AN should support using ANCP to send an Information Report
       message to the NAS indicating the multicast traffic volume that
       has been replicated on that port.

R-69   Upon request by the NAS, the AN must support using ANCP to send
       an Information Report message to the NAS, indicating what
       multicast flows are currently being replicated on a given
       Access Port.

R-70   Upon request by the NAS, the AN must support using ANCP to send
       an Information Report message to the NAS, indicating what
       Access Ports are currently receiving a given multicast flow.

R-71   Upon request by the NAS, the AN must support using ANCP to send
       an Information Report message to the NAS, indicating what
       multicast flows are currently being replicated on each Access
       Port.

R-72   Upon receiving an Admission Response from the NAS, indicating
       that replication of a multicast flow is to start or stop on a
       given access port of the AN, the AN must enforce this decision.
       This decision must be taken irrespective of whether or not a
       corresponding Admission Request was issued by the AN earlier
       on.

4.6.7. Message Handling

R-73   The Access Node must be designed to allow fast completion of
       ANCP operations, in the order of magnitude of tens of
       milliseconds.

R-74   The Access Node should avoid sending bursts of ANCP messages
       related to notification of line attributes or line state, by
       spreading message transmission over time.

4.6.8. Parameter Control

Naturally the Access Node Control Mechanism is not designed to
replace an Element Manager managing the Access Node. There are
parameters in the Access Node, such as the DSL noise margin and DSL
Power Spectral Densities (PSD), which are not allowed to be changed

via ANCP or any other control session, but only via the Element
Manager. This has to be ensured and protected by the Access Node.

When using ANCP for Access Loop Configuration, the EMS needs to
configure on the Access Node which parameters may or may not be
modified using the Access Node Control Mechanism. Furthermore, for
those parameters that may be modified using ANCP, the EMS needs to
specify the default values to be used when an Access Node comes up
after recovery.

R-75   When Access Loop Configuration via ANCP is required, the EMS
       must configure on the Access Node which parameter set(s) may be
       changed/controlled using ANCP.

R-76   Upon receiving an Access Node Control Request message, the
       Access Node must not apply changes to the parameter set(s) that
       have not been enabled by the EMS.

4.7. Network Access Server Requirements

This section lists the requirements for a NAS that supports the use
cases defined in this document. Note that this document does not
intend to impose absolute requirements on network elements.
Therefore, the words "must" and "should" used in this section are not
capitalized.

4.7.1. General Architecture

R-77   The NAS must establish ANCP Adjacencies only with authorized
       ANCP peers.

R-78   The NAS must support the capability to simultaneously run ANCP
       with multiple ANs in a network.

R-79   The NAS must be able to establish an Access Node Control
       Adjacency to a particular partition on an AN and control the
       access loops belonging to such a partition.

R-80   The NAS must support obtaining access loop information (e.g.
       net data rate), from its peer Access Node partitions via the
       Access Node Control Mechanism.

R-81   The NAS must support shaping traffic directed towards a
       particular access loop to not exceed the net data rate learnt
       from the AN via the Access Node Control Mechanism.

R-82 The NAS should support reducing or disabling the shaping limit
      used in the Hierarchical Scheduling process, according to per-
      subscriber authorization data retrieved from a AAA or Policy
      Server.

R-83 The NAS must support reporting of access loop attributes
      learned via the Access Node Control Mechanism to a Policy or
      AAA Server using RADIUS VSAs.

R-84 In a TR-059/TR-101 network architecture, the NAS shapes traffic
      sent to a particular Access Port according to the bitrate
      available on that port. The NAS should take into account the
      layer-1 and layer-2 encapsulation overhead on the Access Port,
      retrieved from the AN via the Access Node Control Mechanism.

R-85 The NAS should support dynamically configuring and re-
      configuring discrete service parameters for access loops that
      are controlled by the NAS. The configurable service parameters
      for access loops could be driven by local configuration on the
      NAS or by a Policy Server.

R-86 The NAS should support triggering an AN via the Access Node
      Control Mechanism to execute local OAM procedures on an access
      loop that is controlled by the NAS. If the NAS supports this
      capability, then the following applies:

      * The NAS must identify the access loop on which OAM
         procedures need to be executed by specifying an Access
         Circuit Identifier (ACI) in the request message to the AN;

      * The NAS should support processing and reporting of the
         remote OAM results learned via the Access Node Control
         Mechanism.

      * As part of the parameters conveyed within the OAM message to
         the AN, the NAS should send the list of test parameters
         pertinent to the OAM procedure. The AN will then execute
         the OAM procedure on the specified access loop according to
         the specified parameters. In case no test parameters are
         conveyed, the AN and NAS must use default and/or
         appropriately computed values.

      * After issuing an OAM request, the NAS will consider the
         request to have failed if no response is received after a
         certain period of time. The timeout value should be either
         the one sent within the OAM message to the AN, or the
         computed timeout value when no parameter was sent.

The exact set of test parameters mentioned above depends on the
particular OAM procedure executed on the access loop. An
example of a set of test parameters is the number of loopbacks
to be performed on the access loop and the timeout value for
the overall test. In this case, and assuming an ATM based
access loop, the default value for the timeout parameter would
be equal to the number of F5 loopbacks to be performed,
multiplied by the F5 loopback timeout (i.e. 5 seconds per the
ITU-T I.610 standard).

   R-87 The NAS must treat PPP or DHCP session state independently from
         any Access Node Control Adjacency state. The NAS must not
         bring down the PPP or DHCP sessions just because the Access
         Node Control Adjacency goes down.

   R-88 The NAS should internally treat Access Node Control traffic in
         a timely and scalable fashion.

   R-89 The NAS should support protection of Access Node Control
         communication to an Access Node in case of line card failure.

4.7.2. Control Channel Attributes

   R-90 The NAS must mark Access Node Control Protocol messages as high
         priority (e.g. appropriately set DSCP, Ethernet priority bits
         or ATM CLP bit) such that the aggregation network between the
         NAS and the AN can prioritize the Access Node Control Protocol
         messages over user traffic in case of congestion.

4.7.3. Capability Negotiation Failure

   R-91 In case the NAS and Access Node cannot agree on a common set of
         capabilities, as part of the ANCP capability negotiation
         procedure, the NAS must report this to network management.

   R-92 The NAS must only commence Access Node Control information
         exchange and state synchronization with the AN when there is a
         non-empty common set of capabilities with that AN.

4.7.4. Adjacency Status Reporting

   R-93 The NAS must support generating an alarm to a management
         station upon loss or malfunctioning of the Access Node Control
         Adjacency with the Access Node.

4.7.5. Identification

R-94 The NAS must support correlating Access Node Control Protocol
      messages pertaining to a given access loop with subscriber
      session(s) over that access loop. This correlation must be
      achieved by either:

      * Matching an Access Circuit Identifier (ACI) inserted by the
         AN in Access Node Control Protocol messages with
         corresponding ACI value received in subscriber signaling
         (e.g. PPPoE and DHCP) messages as inserted by the AN. The
         format of ACI is defined in [TR-101];

      * Matching an ACI inserted by the AN in Access Node Control
         Protocol messages with an ACI value locally configured for a
         static subscriber on the NAS.

4.7.6. Multicast

R-95   The NAS must support using ANCP to configure multicast
       conditional access information to Access Ports on an Access
       Node, using Black Lists, Grey Lists and White Lists.

R-96   The NAS must support using ANCP to indicate to the AN whether
       or not Admission Control is needed for some multicast flows on
       a given Access Port and where needed whether or not the Access
       Node is authorized to perform Admission Control itself (i.e.
       whether or not AN Bandwidth Delegation applies).

R-97   Upon receiving a query from the AN for a request to replicate
       a multicast flow to a particular Access Port, and no AN
       Bandwidth Delegation is used for that flow, the NAS must be
       able to perform the necessary checks (conditional access
       and/or admission control) for the new flow. The NAS must
       support using ANCP to reply to the AN indicating whether the
       request is to be honored or denied. This may involve a
       decision taken locally or querying an external system such as
       a Policy Server.

R-98   Upon receiving a query from the AN for a request to replicate
       a multicast flow to a particular Access Port, and Admission
       Control with AN Bandwidth Delegation is used for that flow,
       the NAS must be able to perform the conditional access checks
       (if needed), and must support using ANCP to delegate a certain
       amount of bandwidth to the AN for a given Access Port.

   R-99   In case of Admission Control with AN Bandwidth Delegation,
          upon receiving a Bandwidth Delegation Request from the AN
          requesting to increase the delegated multicast bandwidth on a
          given Access Port, the NAS must support using ANCP to send a
          Bandwidth Delegation Response indicating the new delegating
          multicast bandwidth.

   R-100 In case of Admission Control with AN Bandwidth Delegation, the
          NAS must support using ANCP to send a request to the AN to
          decrease the amount of multicast bandwidth previously
          delegated on a given Access Port; the NAS should be able to
          specify both the minimum and the preferred amount decrement of
          multicast bandwidth requested.

   R-101 In case of Admission Control with AN Bandwidth Delegation,
          upon receiving an ANCP Bandwidth Release message, the NAS must
          be able to update accordingly its view of the multicast
          bandwidth delegated to the AN.

   R-102 The NAS must support using ANCP to configure the Access Node
          with the "maximum number of multicast streams" allowed to be
          received concurrently per Access Port.

   R-103 The NAS must support using ANCP to incrementally add, remove
          and modify individual entries in White, Black and Grey lists.

   R-104 The NAS must support using ANCP to indicate to the AN whether
          or not multicast accounting is needed for a multicast flow on
          a particular Access Port.

   R-105 In case multicast accounting is needed for a multicast flow on
          a particular Access Port, the NAS should support using ANCP to
          indicate to the AN whether or not additional volume accounting
          information is required.

   R-106 The NAS must support using ANCP to query the AN to obtain
          information on what multicast flows are currently replicated
          on a given Access Port.

   R-107 The NAS must support using ANCP to query the AN to obtain
          information on what Access Ports are currently receiving a
          given multicast flow.

   R-108 The NAS must support using ANCP to query the AN to obtain
          information on what multicast flows are currently replicated
          on each Access Port.

R-109 When Multicast replication occurs on the AN, the NAS must
       support using ANCP to revoke the authorization to replicate a
       multicast flow to a particular Access Port.

R-110 The NAS should support using ANCP to indicate to the AN that
       replication of a multicast flow is to start or stop on a given
       access port of the AN, without having received a corresponding
       Admission Request from the AN earlier on.

4.7.7. Message Handling

R-111 The NAS must be designed to allow fast completion of ANCP
       operations, in the order of magnitude of tens of milliseconds.

R-112 The NAS should protect its resources from misbehaved Access
       Node Control peers by providing a mechanism to dampen
       information related to an Access Node partition.

4.7.8. Wholesale Model

Broadband Forum TR-058 [TR-058], Broadband Forum TR-059 [TR-059] and
Broadband Forum TR-101 [TR-101] describe a DSL broadband access
architecture and how it enables wholesaling. In such a model, the
broadband access provider has a wholesale agreement with one or more
service providers. The access provider owns the broadband access
network and manages connectivity to the service providers. This
allows service providers to provide broadband services to retail
customers, whitout having to own the access network infrastructure
itself.

When applying the Access Node Control Mechanism to a wholesale
network architecture, a number of additional requirements apply.

R-113 In case of wholesale access, the network provider's NAS should
       support reporting of access loop attributes learned from AN
       via the Access Node Control Mechanism (or values derived from
       such attributes), to a retail provider's network gateway
       owning the corresponding subscriber(s).

R-114 In case of L2TP wholesale, the NAS must support a proxy
       architecture that gives different providers conditional access
       to dedicated Access Node Control resources on an Access Node.

R-115 The NAS when acting as an L2TP Access Concentrator (LAC) must
       communicate generic access line related information to the LNS
       in a timely fashion.

R-116 The NAS when acting as a LAC may asynchronously notify the LNS
       of updates to generic access line related information.

5. **Management Related Requirements**

   This section lists the management related requirements for the AN and
   NAS. Note that this document does not intend to impose absolute
   requirements on network elements. Therefore, the words "must" and
   "should" used in this section are not capitalized.

   R-117 It must be possible to configure the following parameters on
          the Access Node and the NAS:

          * Parameters related to the Control Channel transport method:
             these include the VPI/VCI and transport characteristics
             (e.g. VBR-rt or CBR) for ATM networks or the C-VLAN ID and
             S-VLAN ID and p-bit marking for Ethernet networks;

          * Parameters related to the Control Channel itself: these
             include the IP address of the IP interface on the Access
             Node and the NAS

   R-118 When the operational status of the Control Channel is changed
          (up>down, down>up) a linkdown/linkup trap should be sent
          towards the EMS. This requirement applies to both the AN and
          the NAS.

   R-119 The Access Node must provide the possibility using SNMP to
          associate individual DSL lines with specific Access Node
          Control Adjacencies.

   R-120 The Access Node must notify the EMS of configuration changes
          made by the NAS on the AN using ANCP, in a timely manner.

   R-121 The Access Node must provide a mechanism that allows the
          concurrent access on the same resource from several managers
          (EMS via SNMP, NAS via ANCP). Only one manager may perform a
          change at a certain time.

   R-122 The ANCP may provide a notification mechanism to inform the
          NAS about configuration changes done by an EMS, in a timely
          manner. This applies only to changes of parameters which are
          part of the Use Case Access Loop Configuration.

6. Security Considerations

   [draft-ietf-ancp-security-threats] lists the ANCP related security
   threats that could be encountered on the Access Node and the NAS. It
   develops a threat model and identifies requirements for ANCP
   security, aiming to decide which security functions are required at
   the ANCP level.

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

   This is not a new class of risk over already possible IGMP messages
   sent from subscribers to the NAS when the AN uses no IGMP snooping,
   and thus is transparent as long as processing of ANCP messages on the
   NAS/AN is comparable efficient and protected against congestion.

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

   We also observe that an operator can easily deploy some protection
   against attacks using invalid multicast flows by taking advantage of
   the mask-based match in the Black List. This way, joins for invalid
   multicast flows can be denied at the AN level without any ANCP
   protocol interactions and without NAS involvement.

   R-123 The ANCP MUST comply to the security requirements spelled out
          in draft-ietf-ancp-security-threats.

   R-124 The Access Node MUST NOT allow the sending of Access Node
          Control Messages towards the customer premises.

7. IANA Considerations

   This document has no actions for IANA.

8. Acknowledgements

   The authors would like to thank everyone that has provided comments
   or input to this document. In particular, the authors acknowledge
   the work done by the contributors to the Broadband Forum related
   activities: Jerome Moisand, Wojciech Dec, Peter Arberg and Ole
   Helleberg Andersen. The authors also acknowledge the inputs provided
   by Roberta Maglione, Angelo Garofalo, Francois Le Faucheur and
   Toerless Eckert regarding multicast. Finally, the authors thank
   Bharat Joshi, Stefaan De Cnodder, Kirubaharan Dorairaj, Markus
   Freudenberger, Fortune Huang and Lothar Reith for providing comments.

9. References

9.1. Normative References

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

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

[RFC894]   Hornig, C., "A Standard for the Transmission of IP
            Datagrams over Ethernet Networks", RFC 894, April 1984.

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

[draft-ietf-ancp-security-threats]
            Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security
            Threats and Security Requirements for the Access Node
            Control Protocol (ANCP)",
            IETF draft-ietf-ancp-security-threats-06.txt,
            October 2008.

9.2. Informative References

[G.993.2]   ITU-T, "Very high speed digital subscriber line
            transceivers 2 (VDSL2)", ITU-T Rec. G.993.2, Feb 2006.

[G.997.1]   ITU-T, "Physical layer management for digital subscriber
            line (DSL) transceivers", ITU-T Rec. G.997.1, Sep 2005.

[RFC2225]   Laubach, M. and J. Halpern, "Classical IP and ARP over
            ATM", RFC 2225, April 1998.

[RFC2364]   Gross, G., Kaycee, M., Lin, A., Malis, A., and J.
            Stephens, "PPP Over AAL5", RFC 2364, July 1998.

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

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

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

   [RFC4607] Holbrook, H. and B. Cain, "Source-Specific Multicast for
              IP", RFC 4607, August 2006.

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

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

   [TR-147]   Voigt, N., Ooghe, S., and M. Platnic, "Layer 2 Control
              Mechanism For Broadband Multi-Service Architectures",
              Broadband Forum TR-147, November 2008.

Authors' Addresses

   Sven Ooghe
   Alcatel-Lucent
   Copernicuslaan 50
   B-2018 Antwerpen
   Belgium

   Phone: +32 3 240 42 26
   Email: sven.ooghe@alcatel-lucent.com

   Norbert Voigt
   Nokia Siemens Networks
   Siemensallee 1
   17489 Greifswald
   Germany

   Phone: +49 3834 555 771
   Email: norbert.voigt@nsn.com

   Michel Platnic
   ECI Telecom
   30 Hasivim Street
   49517 Petakh Tikva
   Israel

   Phone: + 972 54 33 81 567
   Email: mplatnic@gmail.com

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

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

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

   Phone:
   Email: swadhwa@juniper.net