Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 17, 2007

S. Ooghe
Alcatel-Lucent
N. Voigt
Siemens Networks GmbH & Co. KG
M. Platnic
ECI Telecom
T. Haag
T-Systems
S. Wadhwa
Juniper Networks
February 13, 2007

Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks
draft-ietf-ancp-framework-01.txt

Status of this Memo

Copyright Notice

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 OSS systems.

Table of Contents

1. Introduction

Digital Subscriber Line (DSL) technology is widely deployed for Broadband Access for Next Generation Networks. Several documents like DSL Forum TR-058 [TR-058], DSL Forum TR-059 [TR-059] and DSL 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).

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 NAS), without burdening the OSS layer with unpractical expectations.

Therefore, there is a need 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 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 DSL Forum ([WT-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, 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)

o Line Rate: the total data rate including overhead

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 Messages. The adjacency may either be up or down, depending on the result of the Access Node Control adjacency protocol operation.

o Access Node Control Session: an instantiation of ANCP running on top of the Control Channel. The Access Node Control Session

covers all message exchanges that relate to the actual use cases.

2. General Architecture Aspects

In this section first the concept of the Access Node Control Mechanism is described. Then, the reference architecture is described where the Access Node Control Mechanism is introduced.

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  |---| Aggreg. |---| NAS |---| Network  |
+-----+   +-----+   |        |   |  Node   |   |     |   |          |
| CPE |---| HGW |---|        |   +---------+   |     |   |          |
+-----+   +-----+   +--------+                 +-----+   +----------+
                            Information Reports
                         -------------------------->
                             Control Requests
                         <--------------------------
                             Control Responses
                         -------------------------->

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

Figure 1

From a functional perspective, 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 and/or enforcement function: the reporting function is used to convey status information to the controller function that requires the information for executing local functions. An example of this is the transmission of an Access Loop data rate from an Access Node to a Network Access Server (NAS) tasked with shaping traffic to that rate. The enforcement function can be contacted by the controller function to trigger a local action. An example of this is the initiation of a port testing mechanism on an Access Node.

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 DSL 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 DSL 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 Network Interface Device (NID) at the customer premises, and the Access Node. Legacy protocol encapsulations use multi-protocol encapsulation over AAL5, defined in RFC2684. This covers PPP over Ethernet (PPPoE, defined in RFC2516), bridged IP (IPoE) 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).

2.2.3. Access Node

The Access Node (AN) is a 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). The 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 NAS is a 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. It is often referred to as a Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS). A detailed definition of the NAS is given in RFC2881.

The NAS interfaces to the aggregation network by means of standard ATM or Ethernet interfaces, and towards the regional broadband 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 DSL 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. Access Node Control Mechanism Transport Methods

The connectivity between the Access Node and the NAS may differ depending on the actual layer 2 technology used (ATM or Ethernet). Therefore the identification of unicast & multicast flows/channels will also differ (see also Section 2.4.1).

In case of an ATM access/aggregation network, a typical practice is to send the Access Node Control 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 (e.g. by using a Constant Bitrate (CBR) connection) so that the ATM cells carrying the Access Node Control Messages are not lost in the event of congestion. It is discouraged to route the Access Node Control Messages within the 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 sessions can be routed into a Virtual Path (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 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 Sessions of all Access Nodes. These VLANs should be given a high priority (e.g. by using a high Class of Service (CoS) value) so that the Ethernet frames carrying the Access Node Control Messages are not lost in the event of congestion.

In both cases, the Control Channel between NAS and Access Node can 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 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. In case the response is negative, the state of the peer must then be rolled back to the state prior to the transaction.

2.4. Operation and Management

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.

Also, when using the Access Node Control Mechanism for performing a reporting action, there is a possibility to integrate this with a Subscriber policy system that keeps track of the different Subscriber related parameters.

2.4.1. Circuit 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 or a NAS-Port 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. Dynamic Access Loop Attributes

[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 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 actual data rate, also known as the "DSL sync rate"), hence cannot come from a provisioning and/or inventory management OSS system. Some of the information varies less frequently (e.g. capacity of a DSLAM uplink), but nevertheless needs to be kept strictly in sync between the actual capacity of the uplink and the image the BRAS has of it.

OSS systems are rarely able to enforce in a reliable and scalable manner the consistency of such data, notably across organizational boundaries. The Access Port Discovery function allows the NAS to perform these advanced functions without having to depend on an error-prone & possibly complex 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 sync rate has become stable.

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.

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 (also known as the "DSL sync rate")

o The identification and speed (data rate) of the Remote Terminal(RT)/Access Node link (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 data rate (upstream/downstream)

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

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

- o Minimum 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. If a Subscriber wants to change its Access Loop rate, this requires an OPEX intensive reconfiguration of the Access Port configuration via the network operator, possibly implying a business-to-business transaction between an Internet Service Provider (ISP) and an Access Provider.

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 server (e.g. RADIUS 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, 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 OSS 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 Subscriber 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. This especially holds for an Ethernet-based access/aggregation architecture. In such a architecture, the Access Node, aggregation node(s) and the NAS are involved in the multicast replication process, thereby avoiding that several copies of the same stream are sent within the network.

Typically IGMP is used to control the multicast content replication process within the access/aggregation network. This is achieved by means of IGMP snooping or IGMP proxy in the Access Node, aggregation node(s) and the NAS. However, a Subscriber's policy and configuration for multicast traffic might only be known at the NAS. The Access Node Control Mechanism could be used to exchange the necessary information between the Access Node and the NAS so as to allow the Access Node to perform multicast replication in line with the Subscriber's policy and configuration, and also allow the NAS to follow each Subscriber's multicast group membership.

4. Requirements

4.1. ANCP Functional Requirements

- The ANCP MUST address all use cases described in this document, and be general-purpose and extensible enough to foresee additional use cases (including the use of other Access Nodes than a DSLAM, e.g. a PON Access Node).

- The ANCP must be flexible enough to accommodate the various technologies that can be used in an access network and in the Access Node.

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

- The ANCP MUST be able to recover from loss of ANCP messages.

- 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.

- The ANCP MUST allow fast-paced transactions, in order to provide real time transactions between a NAS an a fully populated Access Node.

- The ANCP MUST allow fast completion of a given operation, in the order of magnitude of tens of milliseconds.

- In large scale networks, Access Nodes are provisioned but not always fully populated. Therefore the ANCP MUST be scalable enough to allow a given NAS to control thousands of Access Nodes (e.g. typically 5000 to 10000).

- The ANCP SHOULD minimize sources of configuration mismatch, help automation of the overall operation of the systems involved (Access Nodes and NAS) and be easy to troubleshoot.

- The implementation of the ANCP in the NAS and Access Nodes MUST be manageable via an element management interface. This MUST allow to retrieve statistics and alarms (e.g. via SNMP) about the operation of the ANCP, as well as initiate OAM operations and retrieve corresponding results.

- o The ANCP SHOULD support a means to handle sending/receiving a large burst of messages efficiently (e.g. using "message bundling").

The ANCP must also support the security requirements as described in Section 7.

4.2. Protocol Design Requirements

- o The ANCP MUST be simple and lightweight enough to allow an implementation on Access Nodes with limited control plane resources (e.g. CPU and memory).

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

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

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

- o The ANCP MUST provide a means for the AN and the NAS to perform capability negotiation and negotiate a common subset.

4.3. Access Node Control Adjacency Requirements

- o The ANCP MUST support an adjacency protocol in order to automatically synchronize states 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.

- o The Access Node Control adjacency MUST be designed such that loss or malfunction of the adjacency can be automatically detected by its peers.

- o The ANCP SHOULD include a "keep-alive" mechanism to automatically detect adjacency loss.

- o A loss of the Access Node Control adjacency MUST NOT affect Subscriber connectivity, nor network element operation.

- o If the Access Node Control adjacency is lost, it MUST NOT lead to undefined states on the network elements.

- o The ANCP MUST be able to recover from loss of the Access Node Control adjacency (e.g. due to link or node failure) and automatically resynchronize state upon re-establishing the Access

Node Control adjacency.

4.4. ANCP Transport Requirements

- o 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.

- o The ANCP MUST be mapped on top of the IP network layer.

- o If the layer 2 transport technology is based on ATM, then the encapsulation MUST be according to RFC2684 routed (IPoA).

- o If the layer 2 transport technology is based on Ethernet, then the encapsulation MUST be according to RFC894 (IPoE).

4.5. Access Node Requirements

This section lists the requirements for an AN that supports the use cases defined in this document.

4.5.1. General Architecture

The Access Node Control Mechanism is defined by a dedicated relation between the Access Node (AN) and the NAS. If one service provider has multiple physical NAS devices which represent one logical device (single edge architecture), then one AN can be connected to more than one NAS. Therefore the physical AN needs to be split in virtual ANs each having its own Access Node Control reporting and/or enforcement function.

- o 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.

- o 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.

- Each AN partition MUST have a separate Access Node Control Session to a NAS and SHOULD be able to enforce access control on the controllers to only designated partitions being bound to one controller.

- The Access Node SHOULD be able to work with redundant controllers.

4.5.2. 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.

- The Control Channel SHOULD use the same facilities as the ones used for the data traffic.

- The Control Channel MUST be terminated at the Access Node.

- For security purposes, the Access Node Control Messages sent over the channel MUST NOT be sent towards the customer premises.

- The Access Node MUST NOT support the capability to configure sending Access Node Control Messages towards the customer premises.

- The Access Node SHOULD process control transactions in a timely fashion.

- The Access Node SHOULD mark Access Node Control Messages with a high priority (e.g. VBR-rt or CBR for ATM cells, p-bit 6 or 7 for Ethernet packets) in order for the packets not to be dropped in case of congestion.

- If ATM interfaces are used, VPI as well as VCI value MUST be configurable in the full range.

- If Ethernet interfaces are used, C-Tag as well as S-Tag MUST be configurable in the full range.

4.5.3. Capability Negotiation

- 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.

4.5.4. Adjacency

- 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.

4.5.5. Identification

- 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.

- To allow for correlation in the NAS, the AN MUST use the same ACI format for identifying the AN and Access Port in Access Node Control Messages, PPPoE and DHCP messages.

4.5.6. Message Handling

- The Access Node SHOULD dampen notifications related to line attributes or line state.

4.5.7. 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.

- 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.

- 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.

4.5.8. Security

The ANCP related security threats that could be encountered on the Access Node are described in [draft-ietf-ancp-security-threats-00.txt]. This document develops a threat model for ANCP security, aiming to decide which security functions are required at the ANCP level.

4.6. Network Access Server Requirements

This section lists the requirements for a NAS that supports the use cases defined in this document.

4.6.1. General Architecture

- The NAS MUST only communicate to authorized Access Node Control peers.

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

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

- The NAS MUST support learning of access loop attributes (e.g. DSL sync rate), from its peer Access Node partitions via the Access Node Control Mechanism.

- The NAS MUST support shaping traffic directed towards a particular access loop to not exceed the DSL sync rate learnt from the AN via the Access Node Control Mechanism.

- The NAS SHOULD support a reduction or disabling of such shaping limit, derived from Policy/Radius per-subscriber authorization data.

- The NAS MUST support reporting of access loop attributes learned via the Access Node Control Mechanism to a Radius server using RADIUS VSAs.

- The NAS MUST correlate Access Node Control information with the RADIUS authorization process and related subscriber data.

- o The NAS SHOULD support shaping traffic directed towards a particular access loop to include layer-1 and layer-2 encapsulation overhead information received for a specific access loop from the AN via the Access Node Control Mechanism.

- o 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 radius/policy server.

- o 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 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).

- o 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.

o The NAS SHOULD internally treat Access Node Control traffic in a timely and scalable fashion.

o The NAS SHOULD support protection of Access Node Control communication to an Access Node in case of line card failure.

4.6.2. Control Channel Attributes

o The NAS MUST mark Access Node Control 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 Messages over user traffic in case of congestion.

4.6.3. Capability Negotiation

o 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.

o 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.

4.6.4. Adjacency

o 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.

4.6.5. Identification

o The NAS MUST support correlating Access Node Control Messages pertaining to a given access loop with subscriber session(s) over that access loop. This correlation MUST be achieved by either:

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

4.6.6. Message Handling

- o 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.

4.6.7. Wholesale Model

- o 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).

- o In case of L2TP wholesale, the NAS MUST support a proxy architecture that enables filtering and conditional access for different providers to dedicated Access Node Control resources on an Access Node.

- o The NAS when acting as a LAC MUST communicate generic access line related information to the LNS in a timely fashion.

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

4.6.8. Security

The ANCP related security threats that could be encountered on the NAS are described in [draft-ietf-ancp-security-threats-00.txt]. This document develops a threat model for ANCP security, aiming to decide which security functions are required at the ANCP level.

5. Policy Server Interaction

This document does not consider the specific details of the communication with a policy server (e.g. using RADIUS).

6. Management Related Requirements

- o 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.

- o 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.

- o The Access Node MUST provide the possibility using SNMP to associate individual DSL lines with specific Access Node Control Sessions.

- o The Access Node MUST notify the EMS of Access Node Control configuration changes in a timely manner.

- o 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.

7. Security Considerations

[draft-ietf-ancp-security-threats-00.txt] investigates the ANCP related security threats that could be encountered on the Access Node and the NAS. It develops a threat model for ANCP security, aiming to decide which security functions are required at the ANCP level. Based on this, the following security requirements are required:

- The ANCP MUST offer authentication of the Access Node to the NAS.

- The integrity of the Access Node Control interactions MUST be ensured using either integrity with a separate protocol (e.g. IPSec) or by designing message integrity into ANCP.

- The ANCP MUST offer authentication of the NAS to the Access Node.

- The ANCP MUST allow authorization to take place at the NAS and the Access Node.

- The ANCP MUST offer replay protection.

- The ANCP MUST provide data origin authentication.

- The ANCP MUST be robust against denial of service attacks.

- The ANCP SHOULD provide mutual authentication between different communicating entities.

- The ANCP SHOULD offer confidentiality protection.

- The ANCP SHOULD distinguish the control messages from the data.

- The ANCP SHOULD provide privacy protection.

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 DSL Forum related activities: Jerome Moisand, Wojciech Dec, Peter Arberg and Ole Helleberg Andersen. The authors also thank Bharat Joshi for commenting on this document.

9. References

9.1. Normative References

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

9.2. Informative References

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

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

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

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

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

[WT-147] Voigt, N., Ooghe, S., and M. Platnic, "Layer 2 Control Mechanism For Broadband Multi-Service Architectures", DSL Forum WT-147, Oct 2006.

[draft-ietf-ancp-security-threats-00.txt]
Moustafa, H., Tschofenig, H., and S. De Cnodder, "Security Threats and Security Requirements for the Access Node Control Protocol (ANCP)", draft-moustafa-ancp-security-threats-00.txt, Dec 2006.

Authors' Addresses

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

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

Norbert Voigt
Siemens Networks GmbH & Co. KG
Siemensallee 1
17489 Greifswald
Germany

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

Michel Platnic
ECI Telecom
30 Hasivim Street
49517 Petakh Tikva
Israel

Phone: + 972 3 926 85 35
Email: michel.platnic@ecitele.com

Thomas Haag
T-Systems
Deutsche Telekom Allee 7
64295 Darmstadt
Germany

Phone: +49 6151 937 5347
Email: thomas.haag@t-systems.com

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

Phone:
Email: swadhwa@juniper.net

Intellectual Property

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

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

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

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).