Network Working Group
Internet-Draft
Expires: November 2, 2006

S. Ooghe
Alcatel
N. Voigt
Siemens
M. Platnic
ECI Telecom
T. Haag
T-Systems
S. Wadhwa
Juniper Networks

May 2006

Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks
draft-ooghe-ancp-framework-00.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 4
1.1. Requirements notation 5
1.2. Definitions 5
2. General Architecture Aspects 6
2.1. Concept of an Access Node Control Mechanism 6
2.2. Reference Architecture 7
2.2.1. Home Gateway 7
2.2.2. Access Loop 8
2.2.3. Access Node 8
2.2.4. Access Node uplink 8
2.2.5. Aggregation Network 8
2.2.6. Network Access Server 9
2.2.7. Regional Network 9
2.3. Access Node Control Mechanism Transport methods 9
2.4. Operation and Management 10
2.4.1. Port Addressing Scheme 11
3. Use Cases for Access Node Control Mechanism 12
3.1. Dynamic Access Loop Attributes 12
3.2. Access Loop Configuration 13
3.3. Remote Connectivity Test 14
3.4. Multicast 15
4. Requirements 16
4.1. ANCP Functional Requirements 16
4.2. Protocol Design Requirements 17
4.3. ANCP transport requirements 17
4.4. Access Node Requirements 18
4.4.1. Access Node Architecture 18
4.4.2. Access Node Control Connection attributes 18
4.4.3. Capability Negotiation 19
4.4.4. Adjacency Requirements 19
4.4.5. Access Node Identification 20
4.4.6. Message Handling 20
4.4.7. Access Node Parameter Control 20
5. Policy Server Interaction 21
6. Security Considerations 22
7. Acknowledgements 23
8. References 24
8.1. Normative References 24
8.2. Informative References 24
Authors' Addresses 25
Intellectual Property and Copyright Statements 27

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. In the scope of these specifications is 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. But due to organizational boundaries between departments operating the local loop, departments operating the ATM network, and departments operating the IP network interworking at the management plane is not always possible. 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.

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

- 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).

- 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 an Broadband Network Gateway (BNG) or Broadband Remote Access Server (BRAS). A detailed definition of the NAS is given in [RFC2881].

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 general-purpose method for multiple network scenarios with an extensible communication scheme, addressing the different use cases that are described throughout this document.

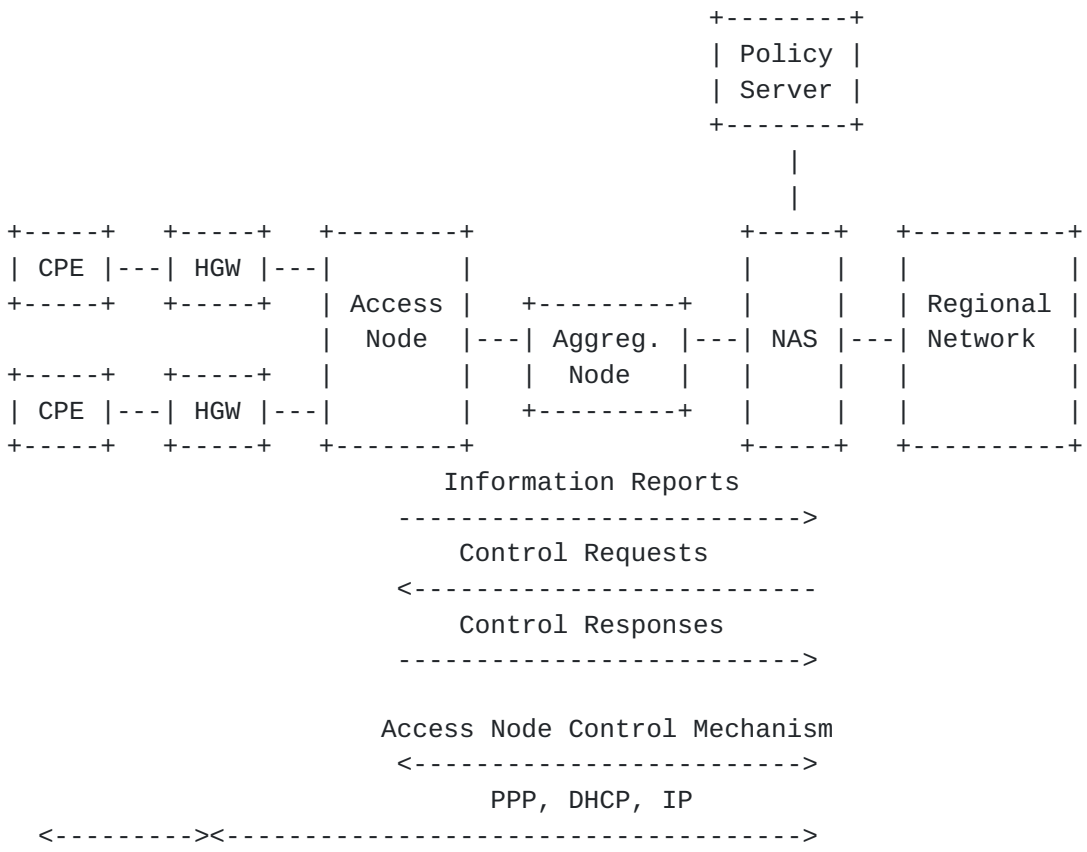

Figure 1

From a functional perspective, a number of functions can be identified:

- 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;

- 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 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:

- 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];

- 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 RFC2364. 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, the Access Node Control messages are sent over a dedicated Permanent Virtual Circuit (PVC) configured between the AN and the NAS. These ATM PVCs should 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 connections 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, the Access Node Control messages are sent over a dedicated Ethernet Virtual LAN (VLAN) using a separate VLAN identifier. Depending on penetration and network design, this can be a simple VLAN for each Access Node or, for higher-grade connections and bundling, one Customer VLAN (C-VLAN) for each Access Node and one Service VLAN (S-VLAN) for all the control connections of every Access Node. 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.

The control connection between NAS and Access Node uses 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 connection.

The control plane interactions are transactional in nature and imply a reliable communication channel to share states. Bidirectional operations are needed, as well as dynamic negotiation of capabilities to address transition issues.

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 central Policy Management system that keeps track of the different Subscriber related parameters (e.g. Access Loop net data rate).

2.4.1. Port Addressing Scheme

In deployments using an ATM aggregation network, the ATM PVC on an Access Loop connects the Subscriber to a NAS. Based on such property, in a PPP scenario, the NAS typically includes a NAS-Port-Id (or NAS-Port in some cases) 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, the port addressing scheme is defined in TR-101. Two mechanisms can be used:

- A first approach is to use a one-to-one VLAN assignment model for all DSL ports. This allows the Access Loop identification to be directly derived from the VLAN tagging, i.e. S-VLAN ID or <S-VLAN ID, C-VLAN ID> pair, of the frames coming from this Access Loop;

- A second approach is to use a many-to-one VLAN assignment model and to encode the Access Loop 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 port 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 the port 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. 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. 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).

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

- The identification and speed (data rate) of the DSL Access Loop (also known as the "DSL sync rate")

- 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]):

- DSL Type (e.g. ADSL1, ADSL2, SDSL, VDSL)
- Framing mode (e.g. ATM, ITU-T Packet Transfer Mode (PTM), IEEE 802.3 Ethernet in the First Mile (EFM))
- DSL port state (e.g. synchronized/showtime, low power, no power/ idle)
- Actual net data rate (upstream/downstream)
- Maximum achievable/attainable data rate (upstream/downstream)
- Minimum data rate configured for the Access Loop (upstream/ downstream)
- Maximum data rate configured for the Access Loop (upstream/ downstream)
- Minimum data rate in low power state configured for the Access Loop (upstream/downstream)
- Maximum achievable interleaving delay (upstream/downstream)
- 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 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.

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 serveral 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 ANCP MUST be an open protocol, either an existing protocol endorsed by an appropriate standard body (e.g. IETF) or a new protocol which will be submitted for standardization to an appropriate standard body. It must be possible for other organizations to define additional protocol information elements.

- The ANCP MUST be transaction-oriented, allowing to reliably share states between the NAS and the Access Node, and recover from loss of synchronization (e.g. node or link failure). Transactions MUST be either fully completed, or rolled-back to the previous state.

- The ANCP MUST be able to recover from access network connectivity disruption and automatically resynchronize. It MUST also be able to recover from message losses on the access network.

- The ANCP MUST allow fast-paced transactions, in the order of magnitude of tens of transactions per second between a given pair (Access Node, NAS). The protocol MUST allow fast completion of a given operation, in the order of magnitude of tens of milliseconds. The protocol MUST be scalable enough to allow a given NAS to control hundreds of Access Nodes.

- 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).

- The ANCP MUST ensure the authenticity of the message initiator and the integrity of the messages. The main goal is to protect the systems (Access Nodes and NAS) against attacks.

- 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.

- A NAS supporting the ANCP MUST correlate layer 2 configuration data with the AAA authorization process and related Subscriber data.

4.2. Protocol Design Requirements

- The ANCP SHOULD have a "boot" sequence allowing to inform the peer about control capabilities supported by the two peers (Access Node, NAS) and negotiate a common subset. This sequence SHOULD be such that a system supporting the ANCP would automatically recognize when its peer doesn't support it at all.

- The ANCP SHOULD include a "keep-alive" mechanism to automatically detect loss of connectivity on the access network or failure of the peer node.

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

- The ANCP SHOULD include a "request/response" transaction-oriented model for the NAS to communicate control decisions or request information from the Access Node. If the response is negative, then the state of the Access Node MUST be unchanged (roll-back).

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

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

- The ANCP MUST be mapped on top of the IP network layer and make use of IPoA or IPoE on ATM or Ethernet, respectively (possibly via a transport layer).

4.3. ANCP transport requirements

- 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.

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

- The transport protocol used for the Access Node Control messages MUST be reliable and scalable.

- A loss of the control connection MUST NOT affect user connectivity and element operation.

- If the connection is lost, it MUST NOT lead to undefined states on the network elements.

- For maintenance purpose the Access Node Control connection MUST be designed in a way that any malfunction of the connection is automatically detected and reported from the Access Node or NAS to the Element Manager or Network Manager to support the network operator in troubleshooting the network.

4.4. Access Node Requirements

4.4.1. Access Node 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) one DSLAM can be connected to more than one NAS. Therefore the physical DSLAM needs to be split in virtual DSLAMs each having its own Access Node Control controler

- 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.

- One partition is grouped of several DSL ports. Each physical DSL port of an Access Node MUST be assigned uniquely to one partition.

- Each AN partition MUST have a separate control connection 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.

4.4.2. Access Node Control Connection attributes

Dependent on network topology the Access Node can be located in street cabinet or central office installation. If an Access Node in street cabinet installation is connected to a NAS all user and Access

Node Control data use the same physical link. Usually, remote Access Nodes are aggregated by an aggregation network and connected to the NAS. Certain connection attributes must be supported

- The Access Node Control Mechanism SHOULD use the same facilities as the ones used for the data traffic.
- The Access Node Control Connection MUST be terminated at the Access Node.
- For security purposes, the Access Node Control messages sent over the control connection 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.
- Control transactions SHOULD be processed in a timely fashion.
- Access Node Control messages SHOULD be marked 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.

4.4.3. Capability Negotiation

- The Access Node MUST provide a capability to indicate which use cases are supported.
- In case the Access Node and NAS cannot agree on a common set of capabilities, the Access Node MUST report this failure to network management.

4.4.4. Adjacency Requirements

- The Access Node MUST be able to support an adjacency protocol used to synchronize states across the link, discover the identity of the entity at the other end of a link, and detect when it changes.
- The Access Node SHOULD report adjacency establishment or loss of adjacency with the controller to network management.

4.4.5. Access Node 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.

4.4.6. Message Handling

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

- The Access Node SHOULD be able to handle sending/receiving a large burst of messages efficiently (e.g. using mechanisms like "message bundling").

4.4.7. Access Node Parameter Control

Naturally the Access Node Control Mechanism is not designed to replace an Element Manager managing the Access Node. There are parameters primarily layer 1 in the Access Node such as the DSL noise margin and DSL Power Spectral Densities (PSD) which are not allowed to be configured by ANCP or any other control connection except the Element Manager. This has to be ensured and protected by the Access Node. It needs to be configured in the Access Node which parameters are allowed (resp. not allowed) to be modified by the Access Node Control Mechanism, how parameters should be modified, stored, or overwritten It needs to be defined the default parameter set the Access Node will come up after recovery.

- It MUST be possible to configure which parameters can be modified.

5. Policy Server Interaction

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

6. Security Considerations

TBD

7. 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 activites: Jerome Moisand, Wojciech Dec, Peter Arberg and Ole Helleberg Andersen.

8. References

8.1. Normative References

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

8.2. 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.

Authors' Addresses

Sven Ooghe
Alcatel
Francis Wellesplein 1
B-2018 Antwerp
Belgium

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

Norbert Voigt
Siemens
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 Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.