Ambient Network Attachment

Teemu Rinta-aho[1], Rui Campos[2], András Méhes[1], Ulrike Meyer[3], Joachim Sachs[1], Göran Selander[1]

Abstract—**The efficiency of network attachment plays a crucial role in the performance of accessing services in new environments. As an example, when a moving network is changing its location relative to attachment points, the detection of the candidate access networks along with their properties and security relationships needs to be carefully managed. This paper presents the framework and mechanisms for network attachment of Ambient Networks. Different steps required for optimizing the network attachment procedure are studied, and a secure network attachment protocol is proposed.**

Index Terms—**Computer network security, internetworking, multiaccess communication.**

I. Introduction

The overall goal of the Ambient Network (AN) Integrated Project [1] is to develop a vision for future wireless and mobile networks. The project aims to create an innovative, industrially exploitable new internetworking framework based on the dynamic composition of networks. A key aspect of the project is to establish a common control layer for various network types, which will provide end users with seamless multi-access connectivity and enable the dynamic selection of the best available network.

This paper outlines the project innovations for network attachment between Ambient Networks. These innovations include cross layer optimizations, early detection of network capabilities and services, guaranteeing service levels across heterogeneous networks and support of symmetric as well as asymmetric attachments. A number of steps are needed for an AN to gain network connectivity in a new location. This may include sensing the media, discovering available networks, selecting the best suitable network(s) and finally negotiating and configuring the network parameters required for data transmission. After this attachment to acquire basic connectivity further negotiations and configurations may take place. In the Ambient Networks project, the network attachment procedure includes only the steps needed to set up a secure communication channel between two attaching ANs .

This paper is organized as follows. In Section II we present our motivation and working assumptions. In Section III we present different options for positioning the network attachment protocol in the layered protocol stack. Section IV discusses the network advertisement and discovery, Section V presents a proposal for the Ambient Network Attachment Protocol (ANAP), and in Section VI possible future work items are presented and the paper is concluded.

II. Motivation and Working assumptions

In this section we provide the motivation for defining a new Network Attachment procedure, so-called Ambient Network Attachment (ANA), and present some working assumptions in defining it.

A. Motivation

Network attachment represents the very first process occurring between two communicating entities (e.g., terminal and network infrastructure). Legacy networking technologies (e.g., WLAN, UMTS) already define the required mechanisms for enabling secure network attachment including address autoconfiguration mechanism that enable global IP connectivity. We may ask why we need to define a new Network Attachment process. One major reason is an attempt to define a generic security handshake independent of access and connectivity, which could potentially reduce attachment time significantly as is explained below. However, there are a number of reasons for enhancing/modifying these legacy attachment procedures. In the following we list the motivations for our work in defining the ANA procedure:

1. **Cross layer optimization:** Experiences with current attachment procedures show a significant inefficiency, in particular due to too loose coupling of different layers. This is most notably the case when multiple independent security associations are established at different layers. While loose coupling provides flexibility, it also results in a number of limitations, including unnecessary latency if the same tasks are performed more than once.
2. **Early and secure detection of network capabilities and services:** Assuming that the available network services depend on the choice of access network, there is a risk that a user needs to perform multiple simultaneous or subsequent attachment procedures over one or multiple access technologies only to find out that the desired network services or capabilities are not available. It is therefore desirable to indicate network capabilities and available services as early in the attachment as possible.
3. **Make service level guarantees across heterogeneous access networks:** From the point of view of heterogeneous access, it is interesting to study if one generic “attachment procedure” (in the sense of security protocol, network service information exchange, etc.) can be applied to any given pair of networks which

[1] Ericsson
[2] INESC Porto
[3] Siemens

have one or more access technologies in common.

4. **Support symmetric and asymmetric settings:** Legacy attachment procedures are typically asymmetric, i.e., they obey to a client-server model, where the roles of each party – a terminal (the client) attaches to an infrastructure network (the server) – and the network services they offer are pre-defined (e.g., terminal runs a DHCP client and the infrastructure network deploys a DHCP server). The communication paradigms assumed for next generation networks and within the AN project consider both asymmetric and symmetric attachments; in the latter, parties have similar capabilities and both can request/offer network services.

B. *Working Assumptions*

In order to clarify the scope and relationship to legacy access technologies, we provide some definitions. The following terms implicitly consider our working assumptions in defining the ANA procedure:

Border Node – a Border Node (BN) of an AN is a physical node sharing a direct physical link with a node outside the AN. Examples of BNs include access points, radio base stations, routers, etc..

ANAP endpoint – an ANAP endpoint is an entity in the ACS (Ambient Control Space) [2] executing ANAP. It is not assumed that the ANAP endpoint is hosted in a Border Node – it could even be several hops away. For example, a BN could be a WLAN AP and the ANAP endpoint is hosted in an AN-enabled WLAN Access Router. In other words, the BN may encapsulate or translate the incoming ANAP messages for transport as required by the internal communication setup of the given AN.

Adjacency – two nodes are adjacent if they are interconnected by a physical link. ANs are adjacent if at least one Border Node in one of the ANs is adjacent to a Border Node in the other AN. The main focus is on attachment between (two) adjacent ANs, but non-adjacent or remote attachment is also considered (see Section V)

III. ANAP AND LAYERED PROTOCOL STACK

To cope with heterogeneous networking environments, ANAP is designed to be independent of the traditional OSI-style layering and any particular interconnecting technologies. Not completely unlike SOAP [5], this independence concerns only the protocol messages; the precise mechanics of binding ANAP to different (combinations of) underlying technologies needs to be defined for each access technology separately. These "bindings" may range from relatively simple adaptations, like a new protocol or payload type definition, to more complicated interventions utilizing reserved fields or otherwise extending existing layers, or even fully integrating ANAP in newly defined technologies.

For a realistic migration story, we assume the co-existence of "AN-aware" and "AN-unaware" layers (i.e. layers with and without appropriate bindings specified), and we allow for "legacy" BNs that do not necessarily (wish to) support these bindings even when they are defined. As a consequence, the actual layer carrying the ANAP messages in any particular attachment instance will depend on which technologies and which bindings the BNs support.

In order to make the best of the different cross-layer optimizations and to secure the exchange as early as possible, ANAP should run on the lowest layer supported by both BNs. Given the mix of legacy, AN-aware, and AN-unaware devices and technologies, much of the complexity for ANAP lies in efficiently finding this lowest common layer without adversely affecting the requisite legacy mechanisms, which often may be running in parallel.

Also in the name of cross-layer optimization, ANAP partially overlaps both "Advertisement and Discovery" (A&D) and "Composition" [9]. The initial roundtrip of ANAP can and often does carry A&D message components, and – in some simple cases – the second (and later) roundtrip(s) may conclude a composition negotiation, e.g. by referencing an already existing composition agreement. While neither of these are strictly part of ANAP per se, the information collected from the various A&D messages on different interfaces and different layers plays an important role in guiding layer selection for ANAP and can form the basis for a subsequent composition negotiation.

We introduce a shorthand notation and some examples of the most relevant options for carrying ANAP messages using different legacy (and in a special case also a non-legacy) technologies.

A. *ANAP in the link layer control plane (L2*)*

In this alternative, the ANAP messages are an integral part of the L2 control plane. This "native ANAP" variant is most suitable for inclusion in new link technologies. Clearly, when such a native mechanism is available, it is – by definition – the lowest common layer for ANAP. As a consequence, failure of ANAP also precludes "legacy" connectivity in this special case.

B. *ANAP in the link layer data plane (L2+)*

Here, ANAP is carried in the data plane of an existing L2. The L2 control plane may or may not be AN-adapted (e.g to include an AN bit or AN information elements in L2 A&D described in the following section). As shown in Figure 1, similar to protocols like ARP or EAPOL, this option would require an "EtherType"-style allocation for ANAP for all supporting L2s.

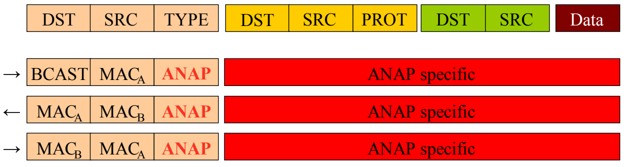

Figure 1: ANAP on L2+

C. *ANAP in the network layer data plane (L3+)*

In this variant, ANAP messages are placed in the payload of existing L3 packets. For layers above L2, the distinction between data plane and control plane seems less relevant. Like for ICMP and GRE, an "IP Protocol"-style identifier for ANAP needs to be allocated on all supporting L3s. Figure 2 shows an example for IPv4 with broadcast and unspecified addresses. Other alternatives include multicast and link-local addresses. Non-IP L3s are expected to afford similar adaptations. Given the widespread adoption of IP and the reluctance to changing existing link layer technologies, ANAP on top of IP appears to be a reasonable short-term option (even for attachment between adjacent networks).

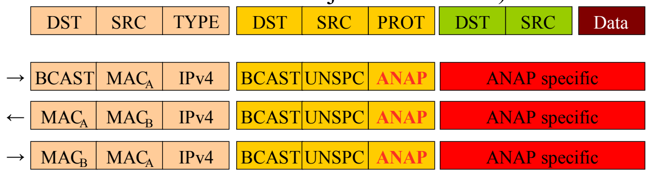

	DST	SRC	TYPE	DST	SRC	PROT	DST	SRC	Data
→	BCAST	MAC$_A$	IPv4	BCAST	UNSPC	ANAP	ANAP specific		
←	MAC$_A$	MAC$_B$	IPv4	BCAST	UNSPC	ANAP	ANAP specific		
→	MAC$_B$	MAC$_A$	IPv4	BCAST	UNSPC	ANAP	ANAP specific		

Figure 2: ANAP on L3+. Shown for IPv4; other L3s similar.

D. ANAP in the transport layer data plane (L4+)

In this case, the ANAP messages are transported in an existing L4. (Naturally, this assumes a working L2+L3 underneath.) Unless the L4 setup procedure is sufficiently advanced to cover dynamic allocation of L4 identifiers, this scheme would necessitate a "TCP/UDP Port"-style allocation for ANAP on all supporting L4s. Since key management protocols usually do not have external reliability requirements, the example in Figure 3 shows UDP on IPv4. Other alternatives abound for both transport and network layers.

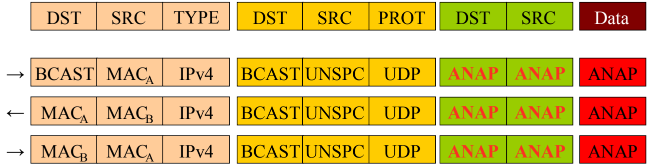

	DST	SRC	TYPE	DST	SRC	PROT	DST	SRC	Data
→	BCAST	MAC$_A$	IPv4	BCAST	UNSPC	UDP	ANAP	ANAP	ANAP
←	MAC$_A$	MAC$_B$	IPv4	BCAST	UNSPC	UDP	ANAP	ANAP	ANAP
→	MAC$_B$	MAC$_A$	IPv4	BCAST	UNSPC	UDP	ANAP	ANAP	ANAP

Figure 3: ANAP on L4+. Shown for UDP on IPv4; other L3s/L4s similar.

While the above list is far from exhaustive, it is thought to cover the most realistic and relevant cases, since even if A&D happens out-of-band (e.g. as an extreme, TV-ads or printed articles may describe AN services), the actual ANAP exchange is likely to follow some combination of the communication patterns listed.

IV. AN ADVERTISEMENTS

A. AN Information Elements

An AN advertisement contains a number of information elements (IE) that carry information necessary for attachment, offered network services or business related information, e.g. tariffs, offered services, etc.

Depending on the underlying access technologies the information elements may be transferred to the receiver of the advertisement by various means, e.g. in beacons, in ANAP messages or by using the mechanisms for inter-AN communication. Information elements should be carefully allocated to various transmission mechanisms in order to balance resource usage and quality of service. Note that this balance may be different for different access technologies. Here are some examples of AN information elements:

- IP configuration parameters (e.g. IP version, IPv6 prefix, DNS server, autoconfiguration mechanism)
- A flag indicating AN support
- Cryptographic IDs of interworking ANs
- Supported QoS classes
- Access type and related service level
- Tariffs, prices of access, payment options

B. AN Advertisements: interworking with legacy access technologies

It is important to understand how existing access technologies could support the distribution of AN Information Elements. Three options are possible.

1) Extended access functionality

AN Information Elements can be included in access layer beacons or other access layer control messages during connectivity setup. This allows any mobile AN to observe the services, capability and characteristics of an access AN at an early stage, before full network attachment is performed. This allows to abort a network attachment procedure at an early stage, if an access AN is found to be insufficient for the desired requirements. This approach requires that access technologies are modified to be able to embed AN Information Elements into access layer beacons or control messages. By broadcasting AN Information Elements, the overhead of the access layer specific beacons increases and leads to a higher signaling cost; at the same time, the effort required by mobile ANs decreases, as they can retrieve significant information by listening to beacons instead of first attaching to a network before receiving further information.

2) Extra bit in access layer beacons

This is a special form of the previous option, however limited to the minimal level of including only a single bit is included in access layer beacons. This AN bit informs the corresponding ANs that the network supports AN technology. Although only little broadcasting overhead is introduced by this bit, it still requires changes to the standards. Although the overhead is small, also the benefit of this option is limited. Knowing that the ANAP procedure could be started gives little indication of the outcome of ANAP.

3) L3 advertisements

This kind of advertisement requires no changes to existing access systems. Before receiving any AN information elements, access layer connectivity is established. After that, AN advertisement can be received via the access layer connectivity. This option has least impact on legacy access system, but it is also least efficient, since connectivity first needs to be established, in order to evaluate in subsequent ANAP and composition procedures, if the connectivity is desired.

V. ANAP PROTOCOL

A. Messages

ANAP defines a family of network attachment protocols. Examples for ANAP are QNAP (Quick NAP) [6] or SNAP (Symmetric NAP) [8]. SNAP is an authenticated Diffie-Hellman key exchange protocol that can also carry other information elements, such as 3rd party authentication protocol messages, address configuration messages etc (cf. Section IV). A detailed description of the SNAP protocol messages and an analysis of the protocol's security properties can be found in [8]. The main goal of SNAP is to provide a symmetric version of QNAP. While QNAP assumes a "provider – subscriber" relationship between the attaching ANs, SNAP does not make such an assumption and is therefore also suitable for network attachments between, e.g., two network providers. Similar to QNAP introduced earlier [6] SNAP also offers piggy-backing of additional information such as IP configuration parameters, the AN-IDs of other composition partners or even QoS capabilities and / or tariffs. As a consequence of the asymmetric setting assumed in QNAP, piggybacking of different information elements onto attachment messages is assumed. SNAP allows for symmetric exchanges of information elements as well. In QNAP the initiator and responder interact with the AAA infrastructure on the responder's side only. SNAP allows each AN to operate an AAA infrastructure with which initiator and responder interact during the protocol execution. SNAP thus supports also network attachments between, e.g., two access networks.

B. ANAP over different access technologies

Depending on the abilities of the underlying access technology the actual use of ANAP differs. With legacy access technologies (e.g. current 3G or WLAN networks) it is not possible to change the existing access attachment procedures. Therefore the legacy advertisements and attachment (incl. security) needs to be done before AN advertising and attachment can be run.

With a modified legacy access technology (e.g. WLAN with extensions) it is possible to send AN advertisements already in the RAT specific beacons. This helps the attaching node in selecting an access network that is an AN, and/or looks otherwise promising service-wise. Otherwise the attachment does not differ too much from the legacy network case. With a future access technology (e.g. Winner [7]), ANAP could be natively supported as *the attachment protocol* of the access technology. In that case access security is not separated from AN security.

C. Remote attachment

It is possible to attach to an AN that is not adjacent, i.e. where no direct links exist between the two networks. We call this remote (or virtual) attachment. In this case, the ANAP protocol runs on layer 3 or higher between the attaching ANs.

To find another AN in the first place, an intermediate network may be used, where remote ANs can be found in e.g. a registration server. The address of the registration server might be well-known to the AN, or it could be discovered in a local registry which is e.g. found through ANAP (comparable to finding a DNS server with DHCP).

Obviously, to get access to the intermediate network in the first place, the ANs have to be attached to a connected network, either using legacy or AN attachment procedures.

As an example, consider a User AN attaching to a Remote Access AN. Firstly, the User AN attaches to a Local Access Network. Then, it discovers the Remote Access AN in an intermediate network registration server (e.g, located in the Local Access Network). During the ANAP handshake a secure bearer is established between the User AN and the Remote Access AN. Using this bearer, additional advertising and composition negotiation can be performed, as if the ANs were adjacent.

D. ANAP and legacy interworking

Previously we have considered attachment over legacy heterogeneous access with the assumption that there is some AN functionality present. In this section we remove all such assumptions and look at the pure legacy case.

Legacy security exchanges - in order to achieve authentication, integrity and ciphering – are usually mandatory. On the other hand, the composition process (for native AN) includes a network attachment which allows running secure authentication and key exchanges or association at the connectivity level.

Therefore, there is a need of defining some new mechanisms allowing an AN to access a legacy network for the further benefits of composition. Among these, it might be that some new business models could be applicable in order to create some additional value for the interworking of the AN and the legacy networks.

The main question is: how to get ANAP messages from AN1 to AN2 if intermediate technologies are AN-unaware. There are two solutions, either to work "on top" of legacy access (i.e. in the IP layer) or to integrate ANAP into the legacy network.

The first solution, namely transporting ANAP messages on or above L2+, leads to three options for connectivity (if supported by legacy access):

1. Use legacy security procedures (in cases where access security is mandatory)
2. Use ANAP to configure L2 security (security is initially completely handled on upper layers e.g. by making ANAP an EAP method, if the legacy network supports EAP [4])
3. Use of legacy security with a special gateway. The gateway then establishes connections to ambient networks

E. After attachment

Prior to the start of the composition agreement negotiation, two ANs establish basic connectivity and security associations by the means of ANAP. With the help of an ANAP protocol two ANs mutually authenticate each other based on their security domain IDs and establish a security association between the ANAP-endpoints in both ANs. This security association may e.g. be an Encapsulating Security Protocol (ESP) security association (SA) as is the case in an ANAP protocol based on HIP [3]. The ANAP messages and messages exchanged after establishing a security association between the ANAP endpoints may be used to exchange advertisement messages (cf. Section IV) that help the ANs to decide whether they want to negotiate a (new) composition agreement. The ANAP protocol may also be used to carry a reference to a pre-established composition agreement. In this case, the composition process can be completed without entering the more complex state of composition agreement negotiation.

The negotiation of a composition agreement may be organized in a centralized or in a decentralized way [9]. In the centralized way one FE (functional entity) in the ACS of each AN negotiates on behalf of all FEs in the AN. In the decentralized way, each FE negotiates with its peer FE directly. Securing the negotiation traffic requires the establishment of security associations either

- between the nodes hosting the negotiating peer FEs, or
- between security gateways through which all negotiation traffic is routed.

It is currently expected that the ANAP-endpoint and the nodes communicating during the composition process do not coincide. As a consequence, not only security associations between ANAP-endpoints but also associations between other nodes will typically have to be established in order to secure the composition process. The required security associations may be:

- Established from scratch
    - E.g. TLS or IKE based on node-IDs
- Derived from the security association established during ANAP
    - by using the security association established during ANAP directly
    - by deriving shared secret key from the ANAP SA and use it on TLS-PSK or IKE-PSK
    - by deriving new SAs directly from ANAP SA

Deriving security associations from the ones established during ANAP seems generally preferable. However, for some types of security associations considered for the protection of the composition related signaling, there are no standardized methods to establish security association between two nodes on behalf of another pair of nodes. This is for example the case for IPsec. Different combinations of endpoints and bootstrapping mechanisms have been described and analyzed in more detail for both types of composition in [9].

VI. Future Work and Conclusion

In this paper we have presented the Ambient Network Attachment procedure, including, the Ambient Network Discovery and Advertisement and an example for the ANAP protocol. In addition, we positioned the ANAP within the layered protocol stack and provided several different options of transporting ANAP between the end-points.

Currently ANAP is implemented in the AN project as two distinct prototypes that consider different aspects of the ANA procedure. The first prototype focuses on the very first attachment sub-process and considers both an extension to the IEEE 802.11 WLAN protocol to support sending of an AN IE in a beacon message as well as a HIP-based attachment protocol. The second prototype focuses on the efficient establishment of IP connectivity between attaching ANs. In the future, we plan to integrate them in a single prototype regarding the implementation of the overall ANA procedure presented herein. Such prototype will afterwards be used for evaluating the ANA procedure in terms of performance and feasibility.

There are still other open issues that need further work. This includes, e.g. clarifying the relationships between the network attachment procedure with other AN FEs. Also, the relation to the composition needs to be clarified – how security associations established between ANAP-enpoints during attachment of ANs can be inherited by other nodes and FEs during the later interactions. Another important aspect is the handover between Border Nodes which belong to the same Security Domain. In that case, it should not be necessary to run full ANAP after handover but instead some kind of a “fast handover” scheme should be used.

Acknowledgment

This paper describes work undertaken in the context of the Ambient Networks project, which is part of the EU’s IST program. Over 30 organizations from Europe, Canada and Australia are involved in this Integrated Project, which runs in 2006-2007 in its second phase. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the Ambient Networks Project.

References

[1] N. Niebert et. al., “Ambient Networks: AN Architecture for Communication Networks Beyond 3G”, IEEE Wireless Communication Magazine, Vol. 11, No 2, April 2004.

[2] Ambient Networks Project: “D1-5: AN Framework Architecture”, IST-2002-507134-AN/WP1-D05, December 2005.

[3] P. Jokela (Ed.), “Host Identity Protocol”, draft-ietf-hip-base-06.txt, June 2006, Work in progress.

[4] H. Levkowetz (Ed.), “Extensible Authentication Protocol (EAP)”, RFC 3748, June 2004.

[5] Simple Object Access Protocol (SOAP), http://www.w3.org/TR/soap/.

[6] J. Arkko, P. Eronen, H. Tschofenig, S. Heikkinen, and A. Prasad, “Quick NAP - Secure and Efficient Network Access Protocol”, Proc.of the 6th International Workshop on Applications and Services in Wireless Networks, May 29-31, 2006.

[7] Göran Klang, “A new radio Interface for the Mobile World – From Vision to Concept: The WINNER Radio Interface Concept”, Wireless World Research Forum, 2005.

[8] Ulrike Meyer, “SNAP: A symmetric version of QNAP”, In Technical Annex to “DG1: Design of Composition Framework”, 2006.

[9] Jorge Andrés-Colás (Ed.), “DG1: Design of Composition Framework”, Ambient Networks Technical Report, 2006.