



WIRELESS WORLD

RESEARCH FORUM

An Overlay Internetworking Architecture for Ambient Networks

Anders Eriksson*, Martin Johnsson*, Andreas Schieder**, Göran Selander*
Ericsson Research

*Torshamnsgatan 23**
SE-164 80 Stockholm
Sweden

*Kackertstrasse 7***
520 72 Aachen
Germany

Abstract—Several different internetworking architectures are being discussed within the Ambient Networks (AN) project. This paper describes the Overlay architecture. Its name is derived from the use of an overlay control plane that allows for internetworking between different network technologies, as well as for control of functionality that is today found in different and sometimes proprietary middleboxes such as NATs, firewalls, mobility agents, and media routers. Routing of control plane signalling messages is based on globally unique, hierarchical, and topologically significant locators. These global locators are mapped one-to-one to local and network technology specific locators such as public and private IPv4 addresses, or IPv6 addresses. Cryptographic host identifiers are used for host identification in line with the Host Identity Protocol framework being developed in the IETF. To reduce the address administration, the routing of signalling messages in access networks can optionally be based on the cryptographic host identifiers. In the user plane, legacy transport protocols such as TCP, UDP, and RTP are used end-to-end. Gateways relay the transport protocol frames between different network technologies, and perform the middlebox functions. The architecture supports backwards compatibility with the legacy IPv4 socket API. Also, migration scenarios where Ambient Networks are interconnected by a legacy IPv4 backbone network are supported.

Index Terms— control plane, internetworking architecture, migration, naming and addressing, signalling

INTRODUCTION

In this paper we propose an internetworking architecture that is designed to support the traditional Internet use cases, as well as future use cases addressed by the Ambient Networks research project. The Ambient Networks use cases include moving networks, media routing, multi-radio resource management, context management, and dynamic set-up of service level specifications [1]. The proposed architecture is based on an internetworking control plane that coordinates the network state for features such as mobility, QoS, and security. The internetworking control plane is designed to control not only IPv4 and IPv6 infrastructure networks, but also networks based on other types of emerging and legacy technologies. A key feature of the architecture is an API that allows for backwards compatibility with IPv4 applications. Also, efficient migration from legacy IPv4 networks is a design goal for the architecture.

The architecture is called the Overlay Internetworking Architecture since it is based on an overlay signaling network.

The Overlay architecture uses several concepts and ideas from previous work in the Ambient Networks project [2], the ITU-T H.323 framework [3], previous work in the area of middlebox control [4][5][6], Split Naming and Forwarding [7], and the Host Identity Protocol (HIP) framework [8]. The

contribution of this paper is the combination and extension of these ideas in the context of an Ambient Network internetworking architecture.

The paper is organized as follows. The high-level requirements on the architecture are described. Based on these requirements, a set of abstract architectural constructs are derived. The concrete architecture is then described in some detail. The paper is concluded with a discussion on the migration characteristics of the proposed architecture.

Problem Statement and High-level Requirements

Considering the diversity of use cases for Ambient Networks, the problem and requirement spaces for the internetworking architecture are quite extensive. The most important problem statements and requirements are highlighted below.

Heterogeneity

The architecture must allow for deployment across virtually any type of network technology, e.g. IPv4, IPv6, MPLS, and for virtually any type of network, e.g. a PAN, a WLAN, a cellular network, or a carrier network. The architecture would then support connectivity across all those networks so that applications perceive transparency regardless of the infrastructure networks that are actually used for connectivity.

Scalability

The architecture shall be possible to instantiate for networks ranging from smallest possible size, e.g. a PAN, up to a large operator network. The architecture shall also provide the means to interconnect the control, user, and management planes of different networks (managed and controlled by different organizations), so that an overlay control, and management plane for all the interconnected networks is provided, which results in a global coherent system for connectivity.

Mobility

The architecture must support that users, applications, devices, and networks can move freely around and attach to different points in the global network. Running applications must not be interrupted due to this mobility, and handovers occurring as a result of the mobility between different points of network attachments shall be transparent to the end users.

Ad-hoc and Moving Networks

Support for ad-hoc networks means capabilities for self-management and self-configuration. There must be support for stand-alone operation, but also for interoperation with other ad-hoc networks in a network island. Finally, network islands should be able to attach to a global network based on the internetworking architecture.

Migration

The internetworking architecture shall offer an API that is backwards compatible with the legacy IPv4 socket API. Moreover, the architecture shall be able to use IPv4 networks as infrastructure networks, as well as backbone networks for the interconnection of AN access networks.

Dynamic Business and Network Environment

The internetworking architecture should be transparent to the business models applied. The architecture must allow for a highly dynamic network environment, and support dynamic network attachments ranging from Personal Area Networks connecting to a residential network, through dynamic interconnection of operator networks.

Security

The AN security architecture requires explicit authorization of resource usage and for that purpose a secure identification scheme. Important objectives are support for user privacy and Denial-of-Service mitigation, which sometimes lead to conflicting requirements on secure identifiers (to be discussed later in this paper).

Architectural Constructs

From the previous section it is possible to analyze and come to an understanding of what concepts, constructs, and traits that are needed and shall be designed into the internetworking control plane architecture. This will then lead to a definition of an architecture in succeeding sections.

Gateways

Gateway functionality is needed to interconnect networks using different technologies, or uncoordinated address spaces, and also for security reasons (e.g. use of proxies).

Overlay

To avoid the need to implement changes in current network technologies, an overlay structure is preferred and should be applied

for all functionality that needs to operate on a global scale and end-to-end. Name-to-address resolution, routing mechanisms, connectivity control, and security are examples where an overlay structure is applicable.

Hierarchy

To address the problems of scalability, a hierarchically ordered structure of resource information, e.g. address ranges, is preferred to enable efficient management, aggregation, and dissemination of such information throughout the global network. Nevertheless, for ad-hoc and moving networks attached at the edge of the global network, or alternatively operating stand-alone, the hierarchical principle could be relaxed without compromising the interoperation with the global network.

Location Transparency

Location transparency refers to the identity of a connection end point, such as a socket, to be independent of where that end point is located in the network topology. This capability allows for session continuity during mobility events of hosts or moving networks. Thus, the definition of a connection end-point should not be based on any name or address, such as a locator, that could change during the lifetime of the connection end-point.

Name Resolution

The internetworking architecture should be able to interoperate with existing name resolution mechanisms, such as DNS, or the GSM HLR and VLR. These name resolution systems would then have to be complemented with a resource record type for AN addressing. Also novel address resolution systems, possibly based on distributed hash tables, should be supported.

Distributed Control

A highly dynamic network environment points directly towards applying distributed control mechanisms, avoiding central functionality. The architecture shall be designed so that the global network is created bottom-up, out from its nodes and its constituent networks.

Secure Identification

There are several aspects of secure identification to be considered in an internetworking architecture. On the one hand we have certain Denial-of-Service attacks that are successful because resources are granted to unknown entities without sufficient control, a security problem that may be addressed with improved

authentication and authorization methods. On the other hand, deployment of a global secure identification schemes such as a Public Key Infrastructure have scalability issues and other problems, in addition to the risks of user privacy violation associated with the traceability and linkability of identifiers.

There are secure identifier proposals that strike a balance between these threats. The features we are interested in come with the use of public keys, or images of functions of public keys, as identifiers, see e.g. Host Identity Tags (HIT) [8] or Keyed Hash Identifiers (KHI) [9]. E.g. in the HIP framework [8] hosts are identified by HITs, a HIT is essentially an image of cryptographic hash functions applied to a public key.

These so called *cryptographic identifiers* have a number of useful features: they come with intrinsic cryptographic properties, such as associated encryption and signature operations; they can be made statistically unique; they can be self-generated or certified directly by trusted third parties or other hosts (in turn represented by cryptographic identifiers); and entities identified with cryptographic identifiers can themselves act both as issuers and subjects of authorisation statements.

Securing identification is necessary for other reasons than host identification e.g. for security management of Ambient Networks and hosts. Also for that purpose, cryptographic identifiers are proposed to identify certain logical entities that execute the management operations.

In order to support the ad-hoc and dynamic interactions featured by AN, the identification scheme must support interaction between strangers while maintaining an appropriate level of security. Additionally to support new and flexible business models for interactions between previously unknown parties, it must be possible to include and securely identify a Trusted Third Party (TTP).

Hence, all AN-enabled nodes are assumed to have cryptographic identifiers of the type mentioned here. By using a suitable authentication protocol, ownership of these identifiers can be cryptographically proven, and some of the denial of service threats mitigated. Self-generated cryptographic identifiers can be used as secure pseudonyms.

The cryptographic identifiers may thus be self-generated, or obtained from/ enrolled with a TTP, depending on purpose and role.

Description of the Overlay Internetworking Architecture Overview and Motivation

The Overlay architecture allows for internetworking over heterogeneous infrastructure network technologies, such as IPv4, IPv6, and other legacy or emerging technologies. It is based on the architectural constructs described in the previous section. In addition to the transport of user data, the architecture supports transport of AN internal signalling from functions such as mobility, QoS, media routing, and context management. The Overlay architecture is thus located at the network layer, and below the session control layer. However, the architecture also supports transport of AN external signalling, such as session control signalling.

The Overlay architecture allows for seamless signaling between Ambient Networks by means of an internetworking control plane with a global and hierarchical address space. The internetworking control plane uses a connectionless overlay signaling network on top of the infrastructure networks, see Figure 1. The infrastructure networks may be heterogeneous in terms of address spaces and network technologies. Gateways are used between such networks to translate network layer addresses and network technology dependent protocols. The gateways are controlled by the internetworking control plane.

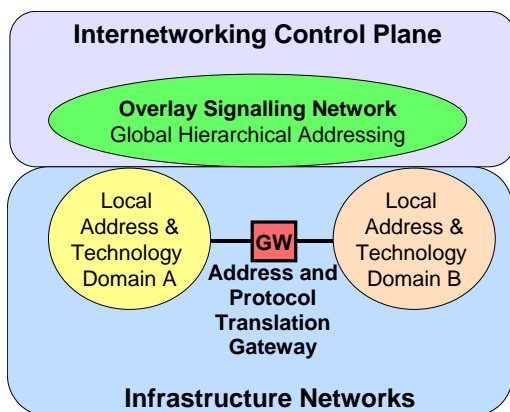


Figure 1: Internetworking principle of the Overlay framework. The gateway is controlled by the internetworking control plane.

Various network technologies with uncoordinated and technology-specific locator name spaces are interconnected by means of the gateways, see Figure 2.

Globally unique and hierarchical addresses in the internetworking control plane are mapped one-to-one to local and network technology specific locators in the infrastructure networks. This mapping is performed by an abstraction layer that adapts the network technology specific control interfaces of the infrastructure networks to a technology-independent control interface, the Ambient Resource Interface [2]. The Ambient Network Interface (ANI) for internetworking, and the Ambient Service Interface (ASI, the AN API), are also shown in Figure 2.

The end-to-end connectivity over the infrastructure networks and the gateways is controlled by means of the network technology independent internetworking control plane that uses the globally unique addresses for routing of signalling messages. In addition to the global and local addresses, each host has a stochastically unique cryptographic identifier to allow for secure identification and to bootstrap other security services such as communication security, secure mobility and multihoming procedures.

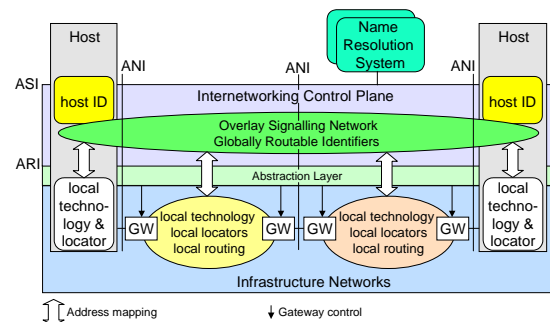


Figure 2: Overview of the Overlay architecture with the overlay signalling network and globally routable identifiers.

The internetworking control plane handles the internetworking between the diverse network technologies. The internetworking is thus based primarily on the control plane, as opposed to the traditional TCP/IP architecture, where a common datagram format and datagram address space is the key to internetworking. However, both the IP and the Overlay architectures use an address space with hierarchical and globally unique addresses. In the Overlay architecture, these global addresses are only used in the internetworking control plane by signaling and routing protocols, while in the IP architecture they are used both in the user plane datagram headers, and in the control plane by the routing protocols as well as by signaling protocols such as ICMP.

Name Spaces

The names and identifiers of the Overlay architecture are illustrated in Figure 3. A Cryptographic Host Identifiers and its 32-bit derivative, the Local Scope Identifiers (LSI), are used according to the host identity and locator split framework [8] to allow for secure identification.

The connection end-points (e.g. sockets) are identified by legacy port numbers plus the LSI. The LSI thus plays the same role in the socket as the traditional IP address. In addition, two layers of locators are used to facilitate routing. This setup adopts the concept of distinguishing clearly between cryptographic host identifiers and locators as mandated by HIP. A major difference is the stack of two different locator planes, which is motivated by the requirement to bridge across independent and unaligned locator domains. The complete setup is depicted in Figure 3 below.

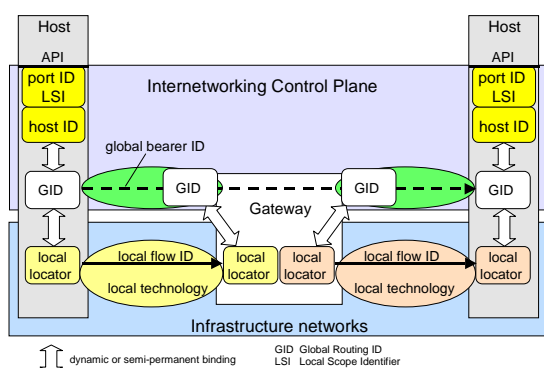


Figure 3: Names and Identifiers used in the Overlay architecture.

The upper locator plane implements the routing overlay, which has given rise to the architecture's name: the Overlay architecture. The locators employed on this plane are denoted Global Routing Identifier (GID). They are assumed to be globally unique and exhibit a hierarchical structure. A GID can represent a network, a node, or an interface depending on the level of detail needed. It can be assigned semi-permanently or dynamically. The primary reason for introducing a hierarchical GID is to allow for scalable global routing. Neither the cryptographic host ID nor the local technology-specific locators are adequate for this purpose.

The GIDs are assumed to be used only for signaling purposes. User-plane data are not assumed to utilize the GID, but continue to use the local technology locators and rely on the gateways to provide a mapping between

the different locator domains. This design choice helps to minimize the impact on the user-plane protocol stacks.

Global and local flow IDs are used to identify flows in the internetworking control plane and in the infrastructure networks respectively. An IPv4 port and address five-tuple is an example of a local flow ID.

Based on the local flow ID of a received packet flow, the gateway uses a look-up table to find the associated global bearer ID, which points at the forwarding context of the flow. The forwarding context specifies the destination local locator to be used in the next-hop network. The forwarding context is set up by the internetworking control plane signalling.

Local technology-specific locators are used in the infrastructure networks. Examples of local locators are IPv4 or IPv6 addresses.

Name Resolution

Several separate external name resolution systems can be employed by the Overlay architecture, such as legacy DNS for the resolution of URIs into GIDs. A new DNS resource record type for GIDs would of course be required. Also, a separate location server system can be used to resolve traditional MS-ISDN numbers to GIDs in the HLR and VLR fashion of GSM.

The external name resolution systems can then be attached to name resolution functions being part of the internetworking architecture, as those external systems mentioned above are not adequate for the handling of highly mobile users and networks. In that respect, the external name resolution systems may store more aggregated address information, which then is further resolved iteratively by the name resolution function in the internetworking architecture until a fully qualified address is obtained.

Routing

When forwarding a signaling message, it must be routed via the appropriate set of gateways between the source and the destination. The routing is at a minimum based on the GID, and different routing schemes are possible to employ in the overlay architecture. A source-directed path like ATM PNNI [11], or alternatively next hop type of mechanisms can be used. In any alternative, when routing takes place of a signaling message towards the destination GID, the hops from gateway to gateway must be done using the local locator technology to

which each of the networks the gateways are connected. This type of address resolution can be based on e.g. an address resolution server per domain using the same method as for IP over non-broadcast multiple access networks [12]. Each gateway is thus configured, possibly dynamically via DHCP [13], with the address of the address resolution server.

The address resolution server caches the GIDs and local locators for each gateway learnt from the received address resolution requests, and uses that information to build an address resolution table for the domain.

An alternative to an address resolution server is to let a dynamic routing protocol, e.g. an extended OSPF, propagate routing information within a domain, and also propagate information about the addresses of the attached gateways [14]. The gateways can then in turn use a BGP style of protocol to disseminate GID routing information among themselves, possibly augmented with a more powerful address and domain aggregation mechanism such as the one employed in ATM PNNI. However, considering the moving network and PAN use cases, these semi-static routing principles must be enhanced with support for mobile hosts, moving networks and ad-hoc networks.

Just like in traditional networks, routing protocols such as BGP and PNNI should work on a longer time scale, while various other mechanisms (e.g. dedicated mobility mechanisms) would handle the dynamics on a shorter time scale.

Internetworking routing information is exchanged between neighbouring gateways using the hierarchical Global Routing Identifier (GID) name space as mentioned above. The topology of the technology-specific connectivity networks is mapped to an abstract technology-independent topology map in the internetworking control plane, see Figure 4. The abstract map only shows the nodes that are controlled by the internetworking control plane, and the abstract networks or links between these nodes. In some cases, an AN may be abstracted as only one node.

In the example shown in Figure 4, the gateways between the different network domains have both technology-specific local addresses and GID addresses with global significance. They are controlled by the internetworking control plane, and are therefore visible in the abstract topology map.

The level of detail in the mapping of the networks between the gateways depends on the required level of resource control. In case of best-effort services, a cloud model of the network between the gateways may be sufficient, while for more predictable QoS levels, control may have to be exercised by the internetworking control plane on a per node and link basis in the network between the gateways. A model for this needs further study, but the ATM PNNI complex node model could be an alternative.

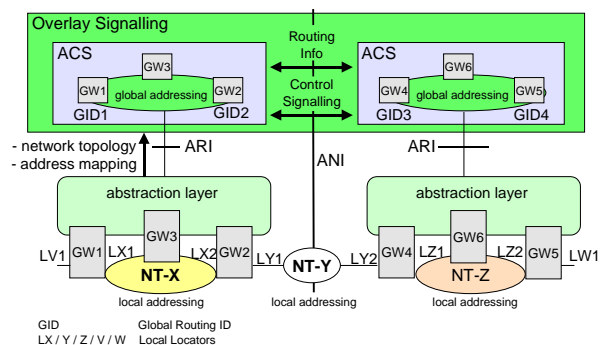


Figure 4: Mapping of the infrastructure topology over the ARI on the technology-independent resource map of the internetworking control plane.

Routing Based on Cryptographic Identifiers

The hierarchical and topologically significant locator (GID) of the Overlay architecture allows for aggregation of locators in the routing tables, which facilitate scalable routing over a global network. However, this incurs off-line administration and assignment of the globally unique locators.

In some access networks of a global network, or in some stand-alone ad-hoc networks, scalable routing may not be an issue. Within such networks, the Overlay architecture supports routing based on the cryptographic host identifier, see Figure 5. A rendezvous server at the border between the global network and the access network registers the host identifiers of the nodes that are present in that network. These nodes also register their locators using a name resolution mechanism that is globally accessible. The locator consists of the GID of the rendezvous server plus the cryptographic host identity.




Figure 5: Routing based on hierarchical locators (GID) in the global network, combined with the option of routing based on cryptographic host identifiers (HID) in some access subnetworks. A rendezvous server (RVS) is used as a gateway between the two types of routing domains.

A signaling message destined to a node in an access network that uses cryptographic host identifiers is routed to the rendezvous server based on the GID for that server. This server then resolves the cryptographic host ID to a local and network technology specific locator, which is used for routing to the host.

Hosts that move between the two types of access networks adapt to the addressing scheme at hand. Where addressing is based on cryptographic identifiers and network technology dependent locators only, the mobile host does not request to be assigned a GID, and does not use the GID layer in the protocol stack.

The use of cryptographic identifiers and rendezvous servers is previously described in the Host Identity Indirection Infrastructure (Hi³) [10]. In particular, the control signaling protocol of Hi³ can be applied to leaf subnets, with certain denial of service attack protection and user privacy benefits.

Note that this section describes the use of cryptographic identifiers for routing only. The use of these identifiers for host authentication purposes is not affected by this routing scheme.

Control and User Plane Protocol Stacks

The Ambient Network Control Space (ACS) is a control plane include functional areas such as QoS and mobility and is distributed on hosts, gateways, and control nodes. It communicates using the connectionless signaling network that is overlaid on the infrastructure networks.

In the user plane, transport protocol frames are carried between the hosts using infrastructure network technologies combined with forwarding contexts that are set up in the

gateways by the internetworking control plane.

The control and user plane protocol stacks for hosts and gateways are shown in Figure 6. Arbitrary infrastructure network technologies L3-X and L3-Y are used as a forwarding layer for control and user data. In the user plane, user transport protocol (UTP) frames are adapted to the network technology at hand using an adaptation layer Ad-X. The user plane transport protocol could be any type of legacy protocol (TCP, UDP, RTP) or future transport protocol. The UTP frames are carried over local network technologies and local address domains. The addresses and technologies used in these local domains are translated in the gateways based on a forwarding context for the bearer that the UTP frames belong to. The forwarding context is set up by the internetworking control plane by means of bearer set-up signaling.

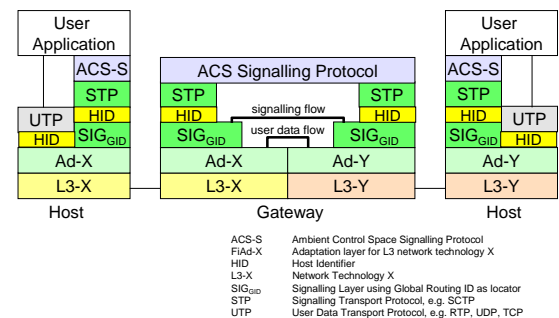


Figure 6: Control and user plane protocol stacks in hosts and gateways.

In the control plane, ACS functional areas communicate based on ACS signaling protocols, which are encapsulated in the Signaling Transport Protocol (STP, e.g. SCTP), see Figure 6. These signaling transport protocol frames are encapsulated in a signaling datagram frame having the GID of the far end host as the destination address (SIG_{GID} in Figure 6). The format of the signaling datagram is for further study. The ACS signaling protocol can be terminated in the gateway. For example, the ACS signaling protocol can request the set up of a forwarding context in the gateway for a specific bearer. There are ACS signalling protocols for functional areas such as QoS, mobility, security, etc.

Alternatively, the gateway can route the signaling datagram based on the GID without terminating the ACS signaling protocol. The routing of signaling datagrams require an intra-AN as well as an inter-AN routing

protocol, in the same fashion as interior and exterior gateway routing protocols in the Internet.

A destination host ID is carried in the ACS signaling protocol. This is used only for authentication and demultiplexing by the hosts, gateways, and control nodes. It is needed for routing only in the access networks that use cryptographic addresses for this purpose.

The transport protocol frames and signaling datagrams are adapted to the various network technologies by means of a network technology specific adaptation layer Ad-X/Y. Gateways between different network technologies relay the transport frames between the technology-specific adaptation layers, i.e. decapsulate the transport frames from one adaptation layer and encapsulate them in another adaptation layer.

Figure 7 shows an instantiation of the protocol stacks with IPv4 and IPv6 as local infrastructure network technologies. In this example, the overlay signaling network uses a subset of the IPv6 address space for the GIDs. The signaling datagram format as such for the encapsulation of the STP frames may or may not be based on IPv6, depending on whether there is a need to reuse IPv6 protocols.

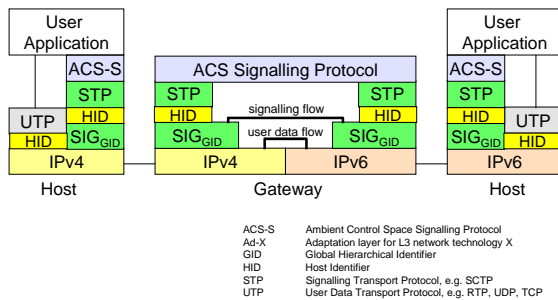


Figure 7: Protocol stack example with IPv4 and IPv6 as infrastructure networks, and IPv6 addressing for the GIDs in the signalling layer.

Figures 6 and 7 do not show the session layer signalling protocol layer. However, a session layer protocol can be encapsulated in a suitable signalling transport protocol and transported over the signalling network between session layer control nodes in the same fashion as ACS signalling protocols. However, the security aspects of allowing the session layer to use the signalling network require further study.

The gateway protocol stacks can be applied also to middleboxes and control nodes.

In principle the control signaling network could be used also for transport of user data. However, this is not a cost-efficient option, since the signaling layer is associated with the overhead of the signaling datagram in Figure 7, e.g. the GID. Moreover, the overlay signaling network would be dimensioned for only a small fraction of the bandwidth of the user plane network.

The protection of the overlay signaling network from denial of service attacks will be investigated in the next phase of the project.

Internetworking Control Plane Signalling

To illustrate the basic operation of the control plane signalling, the set up of an end-to-end bearer is outlined below. However, it should be stressed that the Overlay architecture does not yet prescribe a specific signaling method. The signaling sequence below is included primarily to illustrate the division of functions between the session control network and the Ambient Network, as well as to provide an example of a signaling sequence.

Figure 8 and Figure 9 show the sequence of signals required to set up a basic bidirectional call. SIP/SDP messages are used as examples of session control messages. However, the Overlay architecture is open to any type of session control signalling network that can resolve a high-level destination address into a GID.

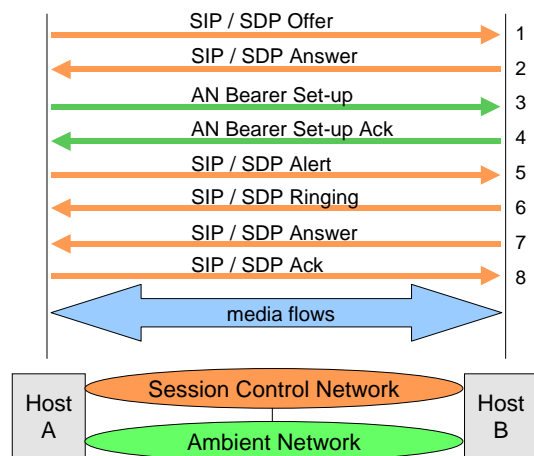


Figure 8: Session and Ambient Network control signalling during call set-up.

The signalling sequence numbering below refers to the numbering in Figures 8 and 9.

1a): The high-level destination user name and service ID are resolved to a host ID, global routing ID (GID) and to a service dependent session protocol such as SIP, with a port number. The division of responsibility between the session control network and the Overlay network for the different steps in the name resolution is for further study. In one scenario, the session control network would request the Overlay network to perform the resolution of the user name into a host ID, and then the host ID into a GID. The Overlay network uses these parameters to forward the session control signalling. They are also used by the session control network when negotiating the session parameters, see step 1b below. The name resolution mechanism could be based on e.g. legacy DNS complemented with location servers. A new DNS resource record type for GIDs that complements the traditional type A record for IPv4 addresses is needed. In an other scenario, the name resolution would be done by traditional SIP proxies.

1b): A SIP/SDP *Offer* message is carried over the overlay signalling network to initiate the session [15]. This message is routed across the Overlays based on the global routing ID (GID). The *Offer* message specifies the session layer parameters that host A would like to use, e.g. the set of media streams, codecs, as well as the GIDs and ports to receive the media streams at host A.

2): Host B responds with an SIP/SDP *Answer* message that specifies the media streams and codecs in the *Offer* message that are acceptable to host B, as well as the GIDs and ports to receive the media streams at host B. This concludes the session layer parameter negotiation.

3a): The originating host uses the GIDs and the port numbers for the media stream destination end-points resulting from step 2 when sending an AN Bearer Set-up signalling Message (ABSM) to request the set up of a bearer. Based on the GID, the ABSM is thus routed to the destination host. This routing is path-coupled, i.e. the ABSM traverses the same gateways as the associated user data bearer. The ABSM can be used to set up unidirectional or bidirectional user data bearers with a specified bandwidth and QoS class. Multicast is for further study.

3b): When the ABSM reaches a gateway, the GID is mapped to a local locator for the next hop gateway. This mapping is based on

routing tables in the gateway that are set up by a combination of interdomain and intradomain routing. This could be done in the traditional Internet fashion, or by means of novel AN routing schemes. The interdomain routing mechanism could be based on the principles of existing internetworking routing protocols, such as BGP or PNNI. The intradomain routing mechanism is a matter of local optimization.

3c): For each gateway that the ABSM traverses, resources for the associated user data bearer are allocated, but not activated. These resources include bandwidth and locators. The resources are allocated in the gateway itself, and optionally also internally in the domain that the gateway belongs to. The mechanism for the domain internal resource reservation depends on the local network technology and is not specified by the Overlay architecture.

3d): When the ABSM reaches the destination end-point, the host ID and the port ID is used to multiplex to the destination socket. The destination application accepts or rejects the bearer specified by the ABSM.

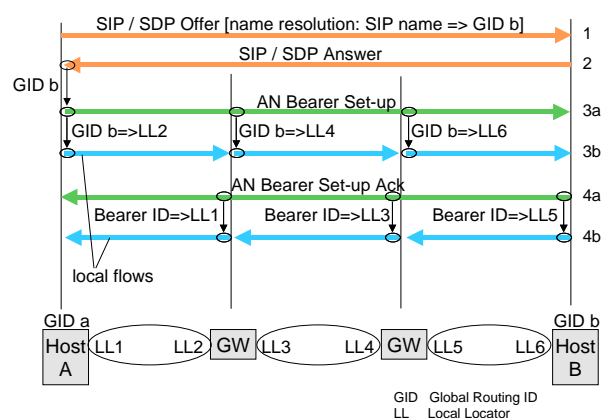


Figure 9: Details of the Ambient Network bearer set-up signalling, steps 1 through 4.

4a): If the bearer is accepted, an ABSM ACK message is routed in the return direction along the same path as the forward ABSM. In each gateway that the ABSM ACK traverses, the previously allocated resources are now activated. In gateways that bridge different address domains, forwarding contexts (locator translation tables, etc.) are set up for mapping of locators for the bearer. Forwarding state per signalling session based on the global bearer ID is maintained in each gateway along the end-to-end path to allow for the ABSM ACK to follow the same path as the forward ABSM. The global bearer ID is used as a look-up key for the forwarding

state specific to the signalling session. (It is for further study whether resources can be activated already by the forward ABSM. This would simplify the procedure, and possibly also make the backtracking of the ABSM ACK unnecessary).

4b): Using the global bearer ID, each gateway looks up the local locator of the previous hop gateway of the forward ABSM to route the ABSM ACK message in the return direction over the same set of gateways as the forward ABSM.

4c): The ABSM ACK message reaches the originating host. The end-to-end bearer is now set up and ready for use by the application.

Figure 8 shows the standard SIP signalling procedure for alert, ringing, answer and acknowledge in steps 5 through 8. However, this session layer signalling sequence is shown only as an example and is not prescribed by the Overlay architecture, which is open to a wide range of session layer protocols.

A legacy IPv4 application may open a TCP or UDP socket directly without the use of SIP signalling. When such an application runs over AN, a traditional DNS query from the host returns a 32-bit LSI that represents the remote host ID, see the HIP framework [8]. The LSI looks like an IPv4 address to the application and is used in place of an IPv4 address when a TCP or UDP socket is opened. The host ID represented by the LSI is then resolved to a GID, which is used when a bi-directional end-to-end bearer is set up according to steps 3 and 4 in Figure 8. This name resolution can be handled by a location server. The name resolution and bearer set-up is handled by the ACS and is invisible to the application.

Since legacy IPv4 applications assume connectionless networks they do not necessarily send explicit bearer tear-down messages. Therefore a time-out mechanism is needed to tear down bearers after a certain period of inactivity. This method is currently used in NATs to tear down translation state.

A more complete signaling solution remains to be worked out as a part of a system specification. The signaling scheme must be optimized based on criteria for mobility, security, QoS, media routing, context management, and multi-radio resource management. This work remains to be finalized. Moreover, considering the dynamic nature of the inter-AN topology, the inter-AN

signalling should probably be based on soft state. Again, the details are a matter for the system specification and are not prescribed by the Overlay architecture.

Note that the total number of signalling messages to set up a call as outlined in Figure 8 could be reduced if the session layer signalling is integrated with the network layer signalling. This is the case for the ISUP signalling of SS7, where four signalling messages perform the same functions as the six signalling messages of steps 3 through 8 in Figure 8.

Note that in a migration scenario where the Overlay internetworking control plane is used over selected subnets of the current best-effort Internet, the signalling scheme described above would be used primarily for the task of setting up NATs and firewalls for a specific end-to-end flow. In this migration scenario the Overlay architecture offers the advantage of a secure multi-purpose signalling mechanism that could replace the current ad-hoc solutions for the control of NATs and firewalls [16] [17]. This is described in more detail in the next section.

Migration

The chances of success for any novel network architecture depend to a great extent on its capability to both add value to, and support migration from, legacy IPv4 networks. Moreover, the novel architecture must support legacy TCP/IP best-effort applications as well as new applications that are able to utilize the enhanced capabilities of an AN. A novel architecture can add value to existing IPv4 networks by providing solutions to problems such as control and management of NATs and other types of middleboxes. It is a design goal for the Overlay architecture to allow for a minimal implementation of the ACS that controls a legacy IPv4 infrastructure network, as well as to allow for backward compatible with legacy IPv4 applications.

In a first migration step, IPv4 and IPv6 access networks are enhanced with an internetworking control plane for control and coordination of functions such as address translation, firewalls, QoS, mobility, and multi-access. This is illustrated in Figure 10. In the user plane, end-to-end transport protocol frames are carried between the edge gateways using standard IPv4 flows over the global IPv4 backbone. The edge gateways control the end-points of these backbone IPv4 flows based on control messages from

the internetworking control plane. Local IPv4 or IPv6 flows over the access networks are concatenated by the gateways with the backbone flows into an end-to-end flow. This is in line with the protocol stack described in Figures 6 and 7. Assuming that the Internet backbone provides adequate QoS, the use of this backbone will be transparent to the ANs attached to its edges.

DNS is used by the ACS to resolve user names to GIDs of destination hosts. Alternatively, DNS resolves the user name to the address of a location server (e.g. a SIP proxy), which stores the GID of a mobile user. The DNS may also store the GID of a rendezvous server at the edge of an access network that uses routing based on cryptographic addresses.

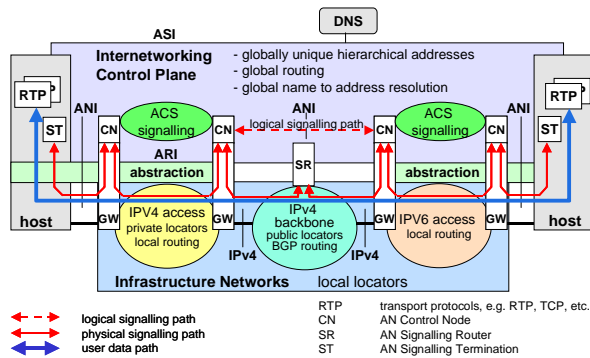


Figure 10: Migration scenario: Communication between Ambient Networks based on IPv4 and IPv6 access network technologies over a legacy IPv4 backbone.

In the internetworking control plane, signalling messages are routed between the gateways across the IPv4 backbone using an overlay network consisting of AN signalling routers. These routers route signalling messages based on the GID. The basic idea is that the signaling router uses a routing table that map GID destination addresses to the IPv4 address of a next hop signaling router, and finally to the IPv4 address of the AN gateway at the edge between the IPv4 backbone and the destination AN. Reachability information is distributed between the signaling routers using policy rules in a fashion similar to BGP. Note that the signaling routers can be implemented as logical nodes in the Overlay edge nodes to avoid the introduction of additional physical nodes in the network.

An example of overlay routing of signalling messages is described in Telephony Routing over IP [19], where telephony routing is performed over the Internet based on an

E.164 number. The TRIP framework describes mechanisms for establishment and maintenance of peering relationships between providers, as well as for exchange and synchronization of gateway routing information between providers. Much of the TRIP framework can be reused in the Overlay signalling router framework.

Alternatively, signalling messages could be routed over the IPv4 backbone simply by retrieving from DNS the IPv4 address of the edge gateway for a specific AN destination network. This approach is in line with legacy routing of SIP messages [18]. However, this method does not provide the same level of policy control of the distribution of the reachability information as the TRIP method.

Note that the applications on both sides of the connection use host identifiers, while the IPv4 or IPv6 addresses are invisible at the application layer, see the protocol stack in Figure 7. Therefore, in the first migration step there is no need for application layer gateways that translate IP addresses used in application layer protocols. Legacy IPv4 applications as well as novel AN applications can communicate seamlessly in this migration step.

Conclusion

We have studied an internetworking architecture based on a combination of globally routable identifiers and cryptographic host identifiers. The two identification schemes complement each other and enable scalable and secure routing.

The description of the Overlay internetworking architecture provided in this paper is sufficiently concrete to serve as a basis for future prototyping activities for the set-up of end-to-end bearers. Such activities should be complemented with a theoretical elaboration of the architecture, primarily in the areas of security and moving networks.

ACKNOWLEDGMENT

We wish to thank András Méhes at Ericsson Research as well as our colleagues in the Ambient Networks project for their valuable contributions to the work that this paper is based on.

REFERENCES

- [1] Ambient Networks deliverable D1.2: AN Scenarios, Requirements, and Draft Concepts.
- [2] Ambient Networks deliverable D1.8: Ambient Networking; Concepts and Architecture.
- [3] ITU-T Recommendation H.323: Packet-based Multimedia Communications Systems.
- [4] IETF middlebox WG:
<http://www.ietf.cnri.reston.va.us/html.charters/midco m-charter.html>
- [5] A. Eriksson, G. Fodor, C-G. Perntz, Middlebox Control Issues in Wireless and Mobile IP Networks. WWRF 11, June 2004.
- [6] A. Eriksson, G. Fodor, A Middlebox Control Plane Framework for Wireless and Mobile IP Networks. INWDA August 2004.
- [7] A. Jonsson, M. Folke and B. Ahlgren, The Split Naming/Forwarding Network Architecture, SNCNETWORK 2003.
- [8] IETF hip WG:
<http://www.ietf.cnri.reston.va.us/html.charters/hip-charter.html>
- [9] A Non-Routable IPv6 Prefix for Keyed Hash Identifiers (KHI), IETF Internet Draft
<http://www.ietf.org/internet-drafts/draft-laganier-ipv6-khi-00.txt>
- [10] P. Nikander, J. Arkko, B. Ohlman: Host Identity Indirection Infrastructure (Hi3)
http://www.ambient-networks.org/docs/Host_Identity_Indirection_Infrastructure_Hi3.pdf
- [11] ATM Forum: Private Network-Network Interface Specification v.1.1
- [12] IETF RFC 2225: Classical IP and ARP over ATM.
- [13] IETF RFC 3646: DNS Configuration options for Dynamic Host Configuration Protocol for IPv6.
- [14] IETF RFC 2328: OSPF version 2.
- [15] IETF RFC 3264: An Offer/Answer Model with the Session Description Protocol (SDP).
- [16] IETF RFC 3489: STUN - Simple Traversal of User Datagram Protocol (UDP) through Network Address Translators (NATs).
- [17] Internet draft: Best Current Practices for NAT Traversal for SIP,
draft-ietf-sipping-nat-scenarios-02 (work in progress).
- [18] IETF RFC 3263: Session Initiation Protocol (SIP): Locating SIP Servers.
- [19] IETF RFC 2871: A Framework for Telephony Routing over IP.