Embedding Identity in Mobile Environments

Alfredo Matos
alfredo.matos@av.it.pt

Susana Sargento
ssargento@det.ua.pt

Rui Aguiar
ruilaa@det.ua.pt

Instituto de Telecomunicações, Universidade de Aveiro
Campus Universitário de Santiago
3800-193 Aveiro, Portugal

ABSTRACT

Recent trends bring Identity concepts into the application layer, although usually focusing in web environments. While this enables new solutions, interactions and paradigms at the application layer, the lower layers are neglected, and considered irrelevant for identity purposes. However, making Identity information available to the OSI stack enables enhanced protocols, which better integrate with A4C mechanisms, and provide better cross-layer integration. We present a solution to integrate identity information into all layers of the OSI stack, and enhance it with resolution mechanisms, enabling full fledged use of Identity by lower layers, such as transport and network. In particular, a new mobility paradigm can be created through an identity-dependent design.

Keywords

Identity Management, Mobility Management, Next Generation Networks.

1. INTRODUCTION

Nowadays, most user services have a notion of identity. This notion takes several forms, such as name and email address or username and password, among many others. With the advent of social internet and the exponential increase of online services, the digital identity of the user becomes a fragmented collection of information bits, scattered across multiple services. To overcome these deficiencies, identity management concepts are appearing at the application level; this helps the user and the network in solving authentication, authorization and privacy problems that derive from the scattered information model.

Currents views on identity tend to isolate the identity layer as something that either has no relationship to the OSI stack, or as something running on top of the application layer. The most common examples of these upcoming technologies are OpenID [12] and Cardspace [13] which leverage identity models to support reliable authentication between the user and its services. With these approaches, a user may choose to use one of several identities that he owns, the one more appropriate to the service at hand; they use identifiers, such as an URL, to point towards user's identity.

But while identity schemes enable new solutions, interactions and paradigms at the application layer, the lower layers are neglected, and considered irrelevant for identity purposes, which are at the core of 4G networks.

Next Generation Networks, or 4G, encompass services ranging from VoIP to IPTV, and span across many different access technologies, such as Wifi, Wimax or even UMTS. Such scenarios are very volatile and user mobility is one of its main characteristics. In such environments, session continuity is very important; therefore, the support of seamless mobility is a main requirement. These networks envision more requirements from those than govern current mobility schemes, such as signal level and network availability. They present managing scenarios that resort to several factors, such as user preferences, provider cost, or QoS preferences. Therefore, mobility must turn to a new place for control. An identity layer, which can accommodate a consistent view over all the necessary parameters, presents the conditions to compel these complex scenarios. Mobility management, driven by identity presents the next step in the evolution of user mobility. Users expect to take full advantage of multiple technologies, without being bound by a specific provider or protocol. In this sense, mobility needs to be addressed in a uniform way, regardless of the mobility protocol used. A common layer should be the enabler to bind different layers, while at the same time allowing them to run independent protocols, such as mobility protocols.

In this paper we propose that identity, building on specific identifiers, should be used to produce an integrated cross-layer design in the network, services and mobility. This identifier needs to integrate at the different layers, enabling a clean identity integration scheme, without requiring modifications to the entire network stack and protocols. This will enable the support of a distributed database model indexed by the same identifier, the provision of an easy usage of user profiles, such as resource authorization and QoS information, and the support of new mobility paradigms. We also present an architecture for , both terminal and the network supporting this identity approach: regardless the addresses used, the same identity material is always provided, greatly simplifying the network processes such as accounting, authorization, QoS reservation. Moreover, using the same identifiers across different attachment points provides a consistent mobility approach across the network: a new mobility paradigm empowered by an identity-dependent design is supported.

The paper is organized as follows: Section 2 discusses the related work. In Section 3 we present the architecture and the necessary modifications to existing paradigms, whereas in Section 4 we show how to integrate Identity Management and mobility. We present a discussion of the integration work in Section 5, concluding and presenting future work in Section 6.

SIGCOMM'07, August 27–31, 2007, Kyoto, Japan.

2. RELATED WORK

While there are no architectures that provide an integration of identity as a design point in a full cross-layer approach, there have been efforts towards this objective. Most of them aim at layer decoupling, providing different namespaces that are or can be used for identity purposes.

The Host Identity Protocol (HIP) [6] is, in its essence, a key exchange protocol based on a new cryptographic namespace. It accomplishes a clean separation of locator and identifier by introducing a new namespace: the Host Identity. A name in the Host Identity namespace, a Host Identifier (HI), represents a statistically global unique name for hosts with an IP stack; it is in fact the identity of a host on the network an it is possible to have multiple identities, some 'well known', some unpublished or 'anonymous'. In order to represent the Host Identity in other protocols, either a Host Identity Tag (HIT) or a Local Scope Identifier (LSI) is used. But, while HIP introduces the concept of an identity namespace, it applies only to a particular layer, while we require a cross layer application of identity.

The Layered Naming Architecture (LNA)[9] aggregates several existing solutions, providing a unified and integrated system. The LNA introduces two layers of names: for service identifiers (SIDs) and for unique endpoint identifiers (EIDs). Both of them are independent of IP addresses. Since LNA is partially based on HIP, it also incorporates the notion of identity, although diluted. The two introduced layers have the objective of providing a decoupled view of different layers, and do not aim at providing a tight integration with identity, even though LNA presents a first step in abstracting sets of identifiers that are not affected by mobility.

The architecture proposed in [1] within the framework of the IST Daidalos project [14], incorporates an authorization model that provides access control to network and services, based on an identity model. It presents security oriented mechanisms for authorization purposes directed at the end user. It also integrates an identifier, ID-Token, that can be seen as a rough identity pointer. Even though it provides some steps towards integration of network authorization with identity models, its main purpose is far from ours: cross layer integration of identity. This architecture can be regarded as a forefather to ours, enhancing network authorization procedures, but cannot be viewed as a complete solution.

Although the aforementioned work provides some steps on the support of a different namespace or identity integration, none is able to provide a unified namespace with a cross-layer design centered on identity and supporting mobile environments.

3. ARCHITECTURE

Integrating identity requires a well defined architecture, which justifies the cross layer integration, enabling next generation mobile environments. To deal with the integration, we also present the refurbishment of the terminal architecture along with the necessary abstractions to render mobility into a paradigm that is leveraged by Identity.

3.1 Network Architecture Model

Typical 4G network scenarios [2],[3],[4], encompass several administrative domains, handled by different providers. They pose a heterogeneous environment powered by different technologies, such as WiFi, Wimax, UMTS or DVB, seamlessly integrated on the architecture. These networks, from where we derive our simplified architecture represented in Figure 1, normally present a controlled environment in terms of resources and authorization. Functional boxes such as Bandwidth Brokers control the network's resources and distribution, facilitating the optimum distribution of resources among the registered users. To sustain the controlled environment, A4C servers take care of terminal authentication and authorization, providing a secure environment for network usage. SIP proxies are in place for the support of SIP based applications, and mobility anchors enable the support of mobility between the several types of networks. A service pool is usually available either local, or remote e.g. Internet. In this 4G network representation we also include and Identity Manager to allow for the support of identity based access to the services and mobility. Sections 3.2 and beyond will further describe this module, its functionalities, and the enhancements it enables in 4G networks.

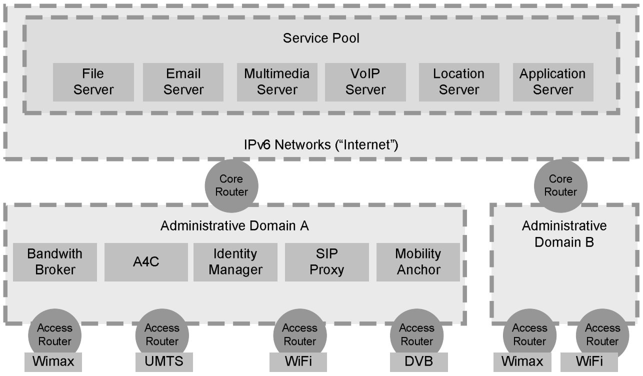

Figure 1. Network Architecture Model

3.1.1 Complexity and Motivation

All of the mentioned functional units have distinct planes of action that culminate in different namespaces for each independent area. But all of them, including services and applications, are user oriented, either for control, management or measurement, leading to different databases on each entity, that in essence deal with information that pertains to the same object – the user.

These properties create conflicts than increase when dealing with mobile environments, such as in the targeted 4G scenarios, where terminal roams across multi-operator heterogeneous networks, using multihoming technologies, which are very controlled in nature in both resources and access. The user has several hurdles to overcome in cooperation with the network, and the first is at the access point, where authentication is required at an A4C, usually provided by PANA [8] or similar protocols. With the generated credentials the A4C will create state at a MIPv6[5] Home Agent, introducing a binding between authentication material and address information. By now several namespaces and identifiers are in play: Layer 2 access between the user terminal and AP; PANA between client, AP and A4C. Possibly Diameter or similar between the A4C and the HA, involving the L3 identifiers.

Afterwards, the terminal will register its L3 addresses with the A4C and the HA, completing the interactions on this plane. But, there are still the interactions with the Bandwidth Broker for QoS purposes. This is established by the terminal, AP or Access router (AR) with the Bandwidth Broker, coordinated with the network entities, using MAC addresses, local IPv6 addresses, Care-of Addresses or Home Addresses.

The next generation mobile environments are expected to provide a user-centric and flexible control over mobility and its granularity, making flow[1] oriented mobility and distribution a necessity. In an optimized scenario, taking into account both user requirements and network resources optimization, flow distribution should depend on interface and network availability, provider information, cost and preferences set by the user while using heterogeneous and disjoint identifiers. Also, new network driven scenarios appear, where the network has control over user mobility [15], clarifying that it is not feasible to stay with the current paradigm, with overhead at several layers and protocols, to enable proper network management relying on volatile and translatable identifiers.

If we take a step forward into multihoming and heterogeneity, then the previous mentioned bindings are multiplied by each active interface, yielding a multidimensional control and data plane where several and different identifiers are used. This complicated environment causes unnecessary and costly mappings, along with unnecessary database replication, where several entries exist across different planes that in fact deal with the same entity.

We propose the usage of cross-layer identifiers is to be used cross-layer in the network and services access and mobility. As previously mentioned, the orthogonal Identity management layer, that currently only exists at the application level services (web and HTTP oriented), can provide the proper space for integration of heterogeneous environments. The namespace provides a rich set of information that directly relates to the user, enabling a set of abstraction that rely on the user-centric paradigms instead of network devices and stack elements. Also, centering the abstraction on the same object – user - enables a consistent cross-layer view, reducing many of the existing ambiguities that occur on the network stack and protocols. Also, different layers of mobility need to be addressed in a uniform way, using the common layer potential to hamper decisions and linkage, regardless of the mobility protocol used, as argued in Section 4. Using the orthogonal layer, applications, transport, network and link layer should be able to correlate services and end-points, reducing the informational overhead and providing a consistent view over user related information. The common layer will be the main support for the binding between different layers, while at the same time allowing them to run independent protocols, including the mobility support.

3.2 Identity Referral

To evolve the aforementioned views on identity onto the network, enabling the same conceptual views to be used across administrative domains for user-centric functions such as QoS or A4C, requires the introduction of two new concepts: an Identity Manager (IDManager) and an Identity Pointer (ID-Pointer). The IDManager stores identity information along with user policies and provides a common view over user information to other network entities, such as domain functions or service providers. It acts upon an identifier that refers to stored identity information. The identifier, ID-Pointer, provides the integration between the Identity Namespace, and consequently the IDManager that represents that namespace, and network protocols. It is used as a handle, derived from identity information, and understandable at the IDManager. For integration purposes (discussed in 3.3) it should also fit the network protocols' identifier space, acting as a cross layer identifier. The structure of the ID-Pointer is shown in **Figure 2**. It is composed of 2 fields: identity *realm* of 16 bits capable of encoding 65536 different identity realms (which can be viewed as an administrative trust domain in Figure 1); and an *Index* of 48 bits capable of indexing 2.8 x 10^{14} different identity registers. Further study should be devoted to achieve a better tradeoff on field sizes. The ID-Pointer will be used throughout the remainder of the paper as a key concept to bring identity to lower layers.

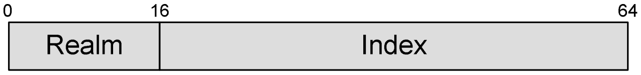

Figure 2. Identity Pointer (ID-Pointer)

The proposed configuration for the ID-Pointer allows any entity on the network to quickly locate the identity's *realm*. Converting a *Realm* into an IDManager's address requires a resolution mechanism. It is out of the scope of this paper to present strategies to resolvable identifiers, but the *realm* can be obtained through a DNS-like mechanism using reverse lookups, or through Distributed Hash Table mechanisms. Through the resolved *realm* an entity has all the necessary information to access the correct IDManager and reach the desired identity information.

Notice that, since the ID-Pointer is a public identifier, it is prone to attacks: the access to the ID-Pointer enables the access to the identity information. Therefore, to gain access to the information, an attacker needs to go through a strong credential/authorization process when accessing the IDManager. A more detailed privacy analysis is done in Section 5.3.

The ID-Pointer is only truly useful if integrated across the network stack. This requires an extension or change on the current protocols and layers, either by using explicit negotiation or by modifying specific layer identifiers to include the ID-Pointer. Along with a sane identity scheme, the cross-layer integration of the ID-Pointer does not require major modifications to the entire network stack and protocols. The next section describes the way this integration is performed.

3.3 Identity Bindings

Upon establishing methods to quickly and easily refer to identities, we need to provide the correct mappings and bindings to be used in a cross-layer design. Using implicit pointers embedded in the protocol, network entities are able to retrieve the identity handle and resolve it without any functional modification to each protocol. Alternatively, identifiers can be exchanged out of band, e.g. using negotiation protocols, requiring network entities to exchange this information deliberately. The optimal solution depends on where and what level we are integrating the ID-Pointer. We present a bottom up approach, covering from the link layer up to the application layer, including mobility, for the

[1] A flow can be defined as an application level association between endpoints, which is an aggregation of source address and port, destination address and port, and protocol.

integration of ID-Pointer structures, either in negotiation phase or imbued in the native protocol identifiers. Figure 3 presents an overview of the ID-Pointer integration at different layers.

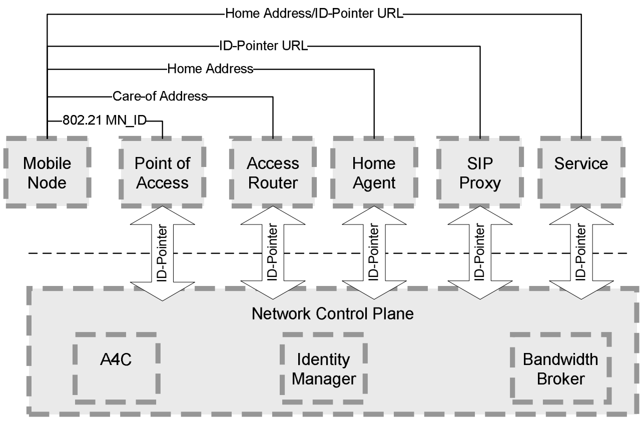

Figure 3. Identity integration at different layers

3.3.1 Link Layer

Since L2 addresses are 48 bit long, there is no space to convey the complete identifier in the addressing structure itself. But, next generation heterogeneous scenarios are using IEEE 802.21 [7] to provide Link Layer Independent Media Services. We can include in the 802.21 negotiation procedures, the ID-Pointer as the MIH Identifier, or in a PANA [8] negotiation phase. Assuming that the reference architecture is 802.21 capable, we replace the MN_ID, which is a TLV field, with the ID-Pointer: this enables the linkage between the MAC address and an MN_ID, therefore providing the reference to the ID Layer.

3.3.2 Network Layer

At layer 3, the IPv6 address provides a proper space to include the ID-Pointer, carried inside the actual locator. The last 64 bits are used to identify the owner of the address, and could be replaced by the ID-Pointer. Figure 4 shows the IPv6 address configuration, built after stateless address auto-configuration, and providing the ID-Pointer to the Access Router (AR). This generated address is in fact the MIPv6[5] Care-of Address (CoA) that will be later registered with the Home Agent (HA). Through the ID-Pointer, the AR has sufficient information to access the network control plane in order to retrieve the required mobile node information, such as QoS, authorization and user preferences.

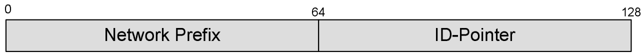

Figure 4. ID-Pointer in the IPv6 address.

3.3.3 Transport Layer

The Home Address (HoA) also follows the same structure, enabling the HA to also have easy access to the mobile node information. Since transport protocols will establish their bindings using the HoA, which acts as the endpoint identifier, we are in fact integrating the ID-Pointer in the transport layer, implicitly conveying identity information also to the services (through HoA).

3.3.4 Application Layer

The application layer has a rich variety of protocols, most of which are URL based. Some protocols that work at the application layer, such as SIP [11], allow the modification of the communication endpoint, while maintaining the ongoing sessions. We base on the SIP example to provide a standard means for integration. Integrating identity into SIP requires breaking the resolution of SIP identifiers into several stages. The terminal registers its HoA with the SIP Proxy, which provides the ID-Pointer. Afterwards, to communicate with a user, one must know the URL, in the form of *johndoe@domain.tld*. The first step is to resolve the *domain.tld* to identify the IDManager, using a DNS record similar to MX, as done for email. The username could then be resolved on the IDManager, obtaining an ID-Pointer, for the target user. This allows re-directing the initiator at the SIP level to the correct resource. This process implicitly links URLs and identities, allowing the initiator to retrieve information from the destination, if allowed, and providing a verifiable identity to the responder - the initiator's ID-Pointer.

3.4 Terminal and Network Support

To have an Identity oriented design, the terminal and network also need to undergo modifications affecting the network control plane.

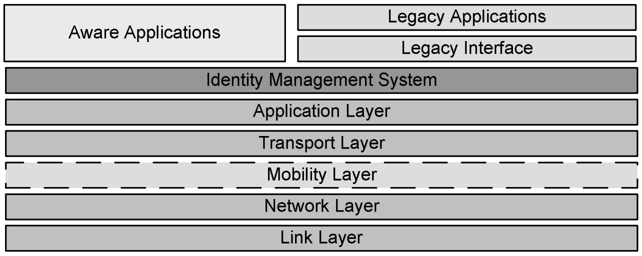

Figure 5. Terminal Control Plane

In the terminal it is required to introduce a control layer that directly instantiates the identity layer functionality, interacting with applications, which provide inputs for network stack management. As seen in Figure 5, applications might be identity aware and provide specific inputs to the management plane, or legacy applications, where the management decisions will be inferred by a legacy interface component. Mobility protocols should refrain providing triggers for mobility, but just reacting to control plane commands, following the identity oriented operations to maintain connectivity. While the control plane has a direct path through the identity management layer, the data plane is orthogonal. Considering that identity management is mainly a control plane task, its repercussion on the data path is to keep each layer consistent with the current identity and mobility policies.

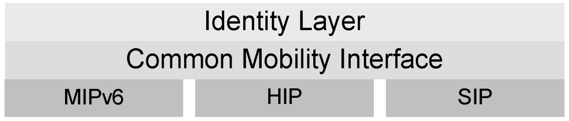

Figure 6. Common interfaces to identity.

In 4G scenarios mobility management cannot be delegated to mobility protocols: with the increase of mobility solutions, it becomes complicated to maintain several sets of policies that apply to particular protocols and situations. We argue that mobility management should be an Identity level function, leaving the mobility protocols as pure signaling mechanisms. With this approach, intelligence is delegated on the same layer, for every protocol, allowing a modular view over mobility signaling, where several protocols can co-exist on the same network stack. The modular approach requires that mobility protocols adhere to the same control interfaces. We introduce the Common Mobility Interface (CMI) that defines a set of primitives that enable the relegation of mobility control and the retrieval of the necessary identity references from the protocol operation. In Figure 6 we present a simple example of the instantiation of the adaptation layer connecting the identity layer with different protocols, which require some extensions, detailed in Section 4.

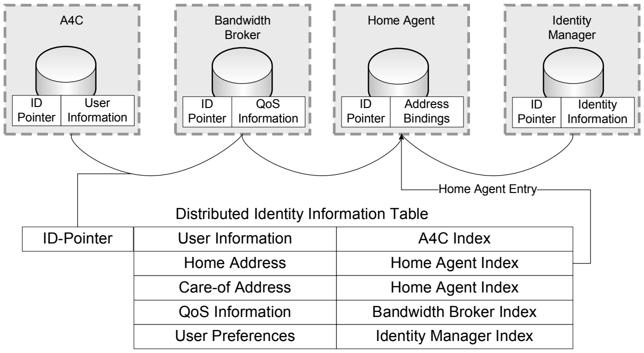

ID-Pointer	User Information	A4C Index
	Home Address	Home Agent Index
	Care-of Address	Home Agent Index
	QoS Information	Bandwidth Broker Index
	User Preferences	Identity Manager Index

Figure 7. Identity Oriented Network Database Model

On the network side the modifications are more operational than functional. The normal network operation is based on distributed information, which can be modeled as relational databases. These databases are unrelated among them, since they use different identifiers for each piece of information, relating to the same user. As shown in Figure 7, we propose to change the way these databases are organized, making them identity oriented, by using the same indexing material, the ID-Pointer, across all the databases: the same ID-Pointer grants access to the relevant database. As an example, using the ID-Pointer at the Home Agent enables the requester to access the HoA and the set of CoAs for a particular database; using it at the Bandwidth Broker grants access to a user's profile.

4. IDENTITY AND MOBILITY MANAGEMENT

In the previous section we introduced the architectural ground stones to accommodate an identity driven architecture that considers mobility as a central piece. Even though the previous generic approaches could be mapped to a wide range of protocols, we selected a few use case protocols to provide clear examples: PANA [8] for authentication, MIPv6 [5] for network layer mobility and NSIS [10] for Quality of Service (QoS). Using these specific protocols, we present generic procedures for bootstrap and seamless handover phases, allowing a clear view of how identity leverages and simplifies protocol mechanisms.

Optimizing user and flow distribution, that governs the mobility process, requires information at several levels. Collecting information about an identity combines the layered view over a user, which is indexed based on ID-Pointers. At the link layer, an 802.21 based framework collects information, such as network availability and provider information, e.g. L2 QoS capabilities. Higher layer information is also easy to retrieve: the ID-Pointer can be used to access QoS information at the Bandwidth Broker, accounting and authorization information at the A4C. User policies can also be involved in the decision process, stored at the IDManager, along with top-level user information. It becomes very easy for the network and user to gather all the necessary information to start the mobility process and run algorithms that effectively distribute the load and optimize resources.

The next sections address the concrete cases of bootstrap and seamless handover leveraged by identity.

4.1 Bootstrap

The bootstrap process consists on a terminal performing the necessary association to enable communication. In Figure 8 we describe the generic approach for network registration and setting up of a flow with a correspondent node. For simplification purposes we suppress the 802.21 signaling assuming that it was correctly performed. After this, the next step is L2 association after which PANA is triggered to perform authentication, where the node conveys the ID-Pointer to the network. The ID-Pointer is then used to contact the A4C that, after resolving the ID-Pointer, contacts the IDManager in order to verify the user identity and obtain the Home Agent for that particular mobile node (to register the necessary MN's HoA). Following these procedures, the A4C notifies the AR that the node identified by the ID-Pointer is authorized with the given HoA. PANA can then be used to convey the HoA to the node, which configures it within its network stack. The terminal is now able to use Stateless Address Auto Configuration (SLAAC) to obtain a CoA and update the Home Agent accordingly. Then, before starting a flow, a QoS reservation is established in the access network: NSIS can also use the ID-Pointer to establish the reservation. All these processes are ID-Pointer based, enabling each network entity to correctly identity the user independently of the protocol and address used.

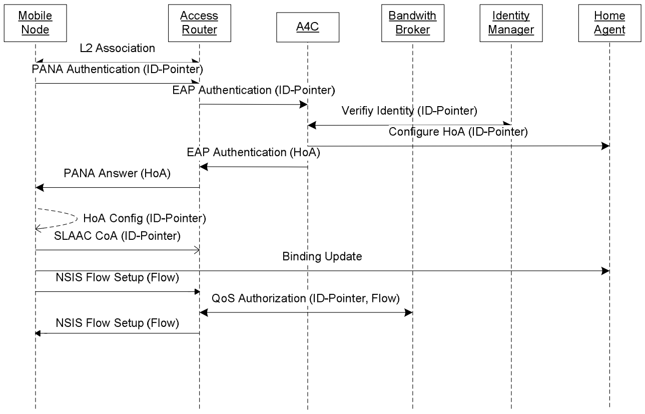

Figure 8. Generalized Terminal Bootstrap Process

4.2 Handover

The trigger to move an identity can have two origins: network or terminal initiated. Identity based mobility has advantages in both cases. When performing terminal initiated mobility, the signaling load can be reduced: the network easily collects the flow information about an identity and performs the movement. In network initiated mobility, the network decides that a particular identity should change the point of attachment. The benefit resides in the ease of information retrieval and signaling, since two access routers can easily share information about an identity, by sharing ID-Pointers and related flow information.

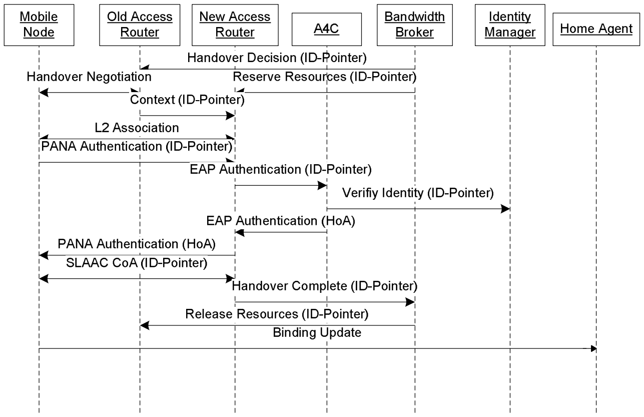
Mobile Node
Old Access Router
New Access Router
A4C
Bandwidth Broker
Identity Manager
Home Agent
Handover Decision (ID-Pointer)
Handover Negotiation
Reserve Resources (ID-Pointer)
Context (ID-Pointer)
L2 Association
PANA Authentication (ID-Pointer)
EAP Authentication (ID-Pointer)
Verifiy Identity (ID-Pointer)
EAP Authentication (HoA)
PANA Authentication (HoA)
SLAAC CoA (ID-Pointer)
Handover Complete (ID-Pointer)
Release Resources (ID-Pointer)
Binding Update

Figure 9. Generalized Network Initiated Handover Procedure

We present a simplified example of a network initiated handover. depicts the information flow process, but with 802.21 signaling suppressed for simplification purposes. A network control point, possibly the Bandwidth Broker, decides to perform the handover of a flow and informs the old and new access routers that the Identity identified by ID-Pointer will change its point of attachment. This is performed by preparing the reservations on the new link, and by the old access router transferring all context relative to ID-Pointer, to the new one. After this process, the node moves and performs a PANA authentication once again (this process could be optimized by pre-authentication schemes before handover). Then, both an address configuration and mobility update are required. In the mobility process, no other elements need to be informed, since the ID-Pointer is still valid and network elements are Identity oriented.

5. DISCUSSION

The discussed identity bindings provide a new integrated view over the network and user environments, while retaining all the protocol properties and assuring backward compatibility. Therefore, we reuse each protocol without modifications, but using the same identifiers across the network, in the form of the ID-Pointer. This section discusses the benefits and costs of such approach, the effect on mobility and support of multiple identities.

5.1 Benefits and Costs

This identity approach allows the support of a distributed database model, indexed at each network element by the same identifier, providing the necessary cross-layer and cross-protocol integration. This distributed meta-system is not bound by a particular protocol identifier, as opposed to today's systems, which require different identifiers at different network points, increasing the problems of the scattered information model.

Moreover, regardless the addresses used, the same identity material is provided, greatly simplifying network processes such as accounting and authorization. Each layer contains information indexed by the same ID-Pointer on both remote and local entities. This means that the index used at the Access Point is the same used at a Bandwidth Broker, and requests and responses are performed based on the same identifier. In complex information environments this enables a uniform view over a user and its sessions, regardless of where they are occurring. In a normal scenario a Bandwidth Broker that keeps track of L2 and L3 assignment would, for the same user, require a list of the MAC addresses in use, a list of IPv6 addresses and their respective Home Addresses, whereas in a common index scenario, only one ID-Pointer would suffice, along with the information necessary in both cases, eliminating several steps of mapping and translation between identifiers, such as Care-of Address and Home Address.

Also, it provides an easier interaction with user profiles. In previous architectures the user profile is not necessarily the user's identity, and in the same reasoning the users identity does not contain network profile information, such as QoS in an Information Card. With the proposed scheme both can coexist on the identity layer, providing information such as resource authorization and QoS information, which are important in a 4G scenario, coupled with the user identity, enabling user-centric architectures. Each entity can retrieve this information easily, whether working with Link Layer or Application Layer information, providing that they have the necessary access credentials to the IDManager.

Current architectures are very intricate and complex systems, making them very hard to innovate since there many design constraints. But, we must oblige by the laid down constraints, such as user centric multi-device environments. Complexity is inherent. The proposed design makes innovation simpler, by reducing the system complexity. Simpler is better in the sense that we can assume simple access to a large set of information, provided by an abstraction layer, valid at each point of the network stack. Network architectures can evolve by turning the focus to truly user-centric paradigms, taking advantage of the provided integration, which fails in current systems.

This also means that the abstraction layer can be summarized in concise APIs, making it simpler to build on top of. Application developers can easily create user-centric software, in a rapid application development environment, since the architecture itself provides the metaphors and handles required by today's business processes.

The cost of such an integrated view is mainly the resolution of the ID-Pointer at each network element: it must be resolved into an IDManager, provided by the *Realm*. This can be done through functions such as reverse DNS queries, or using more evolved mechanisms such as Distributed Hash Tables to locate the correct IDManager. Nonetheless, the cost of resolution can be minimized through caching processes, or optimized through information deduction (e.g. if the A4C receives a preconfigured HoA it can safely infer the *Realm* by looking at the address).

The presented architecture requires a cross-layer introduction of identifiers. While the benefits are several, legacy support of current protocols is a major concern. Taking this legacy requirement into account, we presented integration means that take advantage of the existing identifier spaces in each protocol, bringing to a minimum the impact on each protocol level. The implicit disclosure of the identifiers enables a non-ID enabled node to communicate seamlessly with an ID enabled node, at the cost of neglecting the Identity properties in the communication.

5.2 Effect on Mobility

Mobility has large benefits on the identity approach. Using the same identifiers across different attachment points provides a consistent approach across the network; even when addresses change, the ID-Pointer does not change, and it is only necessary to update the tables directly associated with mobility (since the other non-mobility protocols are bound to the ID-Pointer). It also becomes easier to correlate and coordinate different layer mobility, since the triggers and referrals are consistent.

By shifting the control to an independent layer, it is also possible to modularize mobility, with access to new broader plane of information. From the network point of view, since the identifiers are embedded in the protocols and remain constant, the network is not dependent on the mobility protocol in use for control (e.g. networks do not need to take into account whether route optimization is used with MIPv6, since the ID-Pointer can easily be retrieved in both situations).

5.2.1 Identity-Based Mobility

This paradigm shift has the following immediate impact: instead of moving a particular address to another point of attachment, an identity is being moved. Decoupling the mobility procedures from the terminal as an entity to the identity of the user, increases the overall system flexibility. Since a user is no longer bound to a terminal, an identity can be moved between two different terminals without breaking the mobility mechanisms. Furthermore, with this approach, multiple identities of the same or different users can simultaneously coexist in the same terminal, enabling an identity driven shared environment. This is very useful for mobility purposes, since mobility decisions at the application layer can be linked to identity, with the mentioned benefits.

5.3 Privacy and Multiple Identities

There are two main privacy issues that arise from the presented work: (1) the presence of a resolvable pointer, which enables any passive listener the means to reach identity material; and (2) the integration of identifiers across all layers, raising the possibility of eased linkability of the user's actions.

As previously mentioned, the first issue is solved by a strong security model that requires prior authentication before handing out information to requesters. This implies that prior to the authentication the requester only sees an index and not the identity itself, therefore disclosing no more information than normal protocol operations (although it is the user's choice to disclose a set of public information that might be available). Access to non-public user information is only allowed after the authentication, and even at that point it just conditional and authorized access. Each resolving identity should only access the necessary information, consistent with a minimal information disclosure policy.

As far as linkablility is concerned, if the user deploys the same ID-Pointer across all his network interactions, then there is a high probability of correlating multiple user actions, effectively breaching the user's privacy. However, it is possible to introduce countermeasures that avoid a simple correlation. Using per layer encryption schemes assures that the upper layer identifiers are hidden in transit, and therefore are not linkable by an eavesdropper. This is a scheme which also provides some privacy protection in today's networks, but does not protected from malicious endpoints. A similar protection, but less costly, resides in the capability of using multiple indexes geared at the same the identity, deploying different ID-Pointers for different interactions and layers that point to the same information (several indexes are maintained for the same information), but that is ineffective if the publicly available identity information effectively links the identifiers together. The most effective protection is deploying a multiple identity concept. If different identities can coexist, such as proposed by the high level identity models, then different ID-Pointers can be used for different actions. Passive listeners will not be able to link two actions by resolving two different identifiers, since they do not share information among them. This concept presents the realization of the Information Card metaphor on a network level.

6. CONCLUSION AND FUTURE WORK

This work brings the concept of identity into the whole communication stack, greatly simplifying the network processes such as the support of QoS and authentication/authorization, and presents a vision of mobility that becomes independent of the specific technology protocol. The overall infrastructure that enables this operation requires a distributed linkable database (somewhat implicit already in management systems), changes on resolution systems (to transverse these databases) and on the protocol stack on the equipments. But the most challenging aspect is the need for a new vision in the mobility world – now re-centring itself around the identity of the customer in all its actions. This opens the path to a decoupling of mobility management (user, device, session) from underlying technologies, smoothing network evolution and driving optimization aspects at all levels of the OSI stack.

Future work to be performed is the mapping of this architecture in specific protocol instances, namely those that seem to provide a simpler evolution path from current 3G networks, and to evaluate its performance and scalability. The underlying idea is to be able to provide a paradigm change on network operation, embedding (customer) identity in the regular network management and control. This will lead to a mobility control which will increasingly become independent from the underlying technology, providing easy migration paths, with evolving mobility-aware services.

7. Acknowledgements

We would like to thank João Girão, NEC Network Labs, for his fruitful discussions.

The work described in this paper is based on results of IST FP6 Integrated Project DAIDALOS. DAIDALOS receives research funding the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this paper.

8. REFERENCES

[1] Olivereau, A. et all, "An Advanced Authorization Framework for IP-based B3G Systems". In *Proceedings of the 14th IST Mobile & Wireless Communications Summit*; Dresden, Germany, June 19-23, 2005.

[2] Julien Abeillé, et al, "MobiSplit: a scalable approach to emerging mobility networks", *First International Workshop on Mobility in the Evolving Internet Architecture* (MobiArch 2006), pp 17-22, Dec 1, 2006.

[3] Aguiar, R., Einsiedler, H., Karrer, R., "Daidalos: The Operators Vision of the Next Generation Internet", in *Infocom 2006,* Barcelona, Spain, April 23-29.

[4] ITU-T Y.2011, “Next Generation Networks – Frameworks and functional architecture models, General principles and general reference model for Next Generation Networks”, 10/2004

[5] Perkins, C. et al, “Mobility Support in IPv6 (MIPv6)”, IETF RFC, June 2004.

[6] R. Moskowitz, “Host Identity Protocol” IETF Internet Draft (Work in Progress), July 2006.

[7] IEEE P802.21/D04.00, “Draft IEEE Standard for Local and Metropolitan Area Networks: Media Independent Handover Services”, March 2007

[8] Frosberg, D et al, “Protocol for Carrying Authentication for Network Acess (PANA)”, IETF Internet Draft (Work in Progress), March, 2007.

[9] H. Balakrishnan, et al. “A Layered Naming Architecture for the Internet”. *SIGCOMM*, Portland OR, Aug 2004.

[10] Xiaoming Fu et al., “NSIS: a new extensible IP signaling protocol suite”, IEEE Communications Magazine, October 2005:133-141

[11] J. Rosenberg, et al, “SIP: Session Initiation Protocol”, IETF RFC 3261, June 2002.

[12] D. Recordon and B. Fitzpatrick, “OpenID Authentication 1.1”, May 2006.

[13] Windows Cardspace, Internet Address: http://cardspace.netfx3.com/content/introduction.aspx

[14] IST FP6 Integrated Project Daidalos: http://www.ist-daidalos.org

[15] Aguiar, R, Banchs, A, Melia, T, Pacyna, P, “NIHO: Network Initiated Handovers for next generation ALL IP Networks”, in *Infocom 2006*, Barcelona, Spain, April 23-29, 2006.