Privacy and Identity Management in a Layered Pervasive Service Platform

Marc BARISCH[1], Martin NEUBAUER[1], Joao PAGAIME[2], Joao GIRAO[2], Rui L. AGUIAR[3]
[1]University of Stuttgart, Institute of Communication Networks and Computer Engineering,
Pfaffenwaldring 47, 70569 Stuttgart, Germany
Tel: +49 (711)685-60217, Fax: +49(711) 685-50217,
{barisch,neubauer}@ikr.uni-stuttgart.de
[2]NEC Europe Ltd., Kurfürsten-Anlage 36, 69115 Heidelberg, Germany
Tel: +49 (6221) 4342-0, Fax: +49 (6221) 4342-155
{joao.dasilva,joao.girao}@nw.neclab.eu
[3] Instituto de Telecomunicaes/Universidade de Aveiro, 3810-193 Aveiro, Portugal
Tel: +351 234 377900, Fax: +351 234 377900
ruilaa@det.ua.pt

Abstract: Making pervasive computing reality is a challenging task mainly due to the multitude of functional requirements and technological constraints. In parallel to the honourable research progress in specific technologies, the Daidalos project assessed that in future there will be the need for a pervasive service platform with open interfaces in order to simplify service development and provisioning. The success of such a platform depends on the balance of different aspects, e.g. operational costs with revenue potentials, collection of personal data for context-awareness with privacy protection, manual control and transparency with enhanced user experience and simplicity. In this paper we show the Daidalos approach to privacy protection and identity management for a future pervasive service platform and its architecture. We show how user identities are structured to support dynamic context information while following regulations for privacy protection in Europe. Special focus is put on the trade-off between access control for privacy protection and user experience. This is achieved by automated identity selection, automatic derivation of fine-grained access control policies and their deployment. We also present gathered performance data and implementation details of our ID Broker concept.

Keywords: Identity Management, Privacy Protection, Virtual Identity

1. Introduction

Academia and industry have been working towards Weiser's vision [1] of future computing. This vision materialises in recent advances in wireless communication and miniaturisation of technologies. These allow provisioning of computing power, communication, storage and sensing of the surroundings every time and everywhere. Consequently, systems become more dynamic, e.g. interaction partners change more often and interactions themselves adapt to changes in context. Context-awareness implies increased collection and processing of personal data, which amplifies the potential to invade users' privacy. To mitigate this flaw, solutions must be integrated to protect personal data while "they weave themselves into the fabric of everyday life until they are indistinguishable from it" [1].

The Daidalos project shares this ubiquitous computing vision and is designing a layered service platform with inherent privacy protection. The general approach to privacy protection is to guarantee the right on informational self-determination and allow users to act pseudonymously in the platform by means of *virtual identities* (VIDs). Different VIDs can

be used to consume services. The view on user attributes associated to an identity is limited. Hence, VIDs and identity management (IdM) are key concepts in the enforcement of privacy protection as well as for success and acceptability of pervasive services. This work is in accordance with legislative regulations in Europe [2, 3].

Existing approaches do not fulfil all requirements of pervasive service platforms with respect to privacy. In particular, they mostly regard privacy protection as pure manual control over which attributes are disclosed to which service on application layer. They neither provide mechanisms for dynamic access control on attributes nor do they give recommendations whether it is sensible to disclose the attributes in this context. If disclosure of attributes is not restricted, a malicious service provider could gain a detailed view on an identity.

In addition, if multiple VIDs per entity are supported then it has to be ensured that no entity can link them by evaluation of information inherent to the system. For example, using different VIDs for different services that are provided by a single service provider concurrently are easily linkable in today's systems by evaluating the IP-address. To overcome this flaw, cross-layer considerations and solutions such as [4] should be taken into account. To the best of our knowledge, existing approaches do not consider linkability of VIDs previously used in different contexts and across layers.

Eventually, performance of retrieving personal data matters in case of frequent changing values like location or sensor data. This means, the mobile terminal should not be involved in the retrieval due to the scarce radio resources.

In Section 2, we first give a brief overview of related work in the IdM area. Section 3 gives an example scenario, which can be realised by the Daidalos architecture. The underlying identity model is presented in section 4. Section 5 and 6 take a closer look at the components of the proposed IdM architecture. Section 7 summarises and concludes our work.

2. Related Work

The specifications of Shibboleth [5] and Liberty Alliance [6] provide federated identity management with main focus on Single Sign-On and Logout for Web Services. Both are based on SAML [7] but address different target groups. Shibboleth concentrates on academia and provides mechanisms for direct exchange of user attributes between federated institutions for authorisation purposes. Institutions and users can protect their attributes by so called attribute release policies, which can be modified by graphical user interfaces [8]. In contrast, Liberty Alliance concentrates on business environments and has a set of services based on templates [9], which allow the management and release of user attributes. The concrete realisation of attribute release policies is not specified.

Microsoft CardSpace [10] uses WS-* specifications and requests the user to select an identity, which is termed identity card. Each identity card comprises a set of attributes, so called claims. These claims are transferred from the card provider (=identity provider) to the service provider via the user's terminal. This means that there is no direct communication between the identity provider and service provider, which contrasts the approaches of Shibboleth and Liberty Alliance.

Beyond the web service world, approaches for attribute management are specified by OMA and 3GPP. OMA has specified the PEEM [11], which is a generic policy-based approach for the control of service delivery platforms.

3GPP specified the "Generic User Profile (GUP)" [12]. It allows services to access user data by contacting a central point and thus has the possibility to hide several data stores. The GUP concept is closely related to our approach. However, we focus on pervasive computing and privacy protection.

3. Daidalos Architecture and Usage Scenario

Daidalos is a FP6-funded European research project with the goal to create a service platform, which enables pervasive services for mobile users with a major focus on privacy protection. In addition the platform has to ensure that telecommunication operators as well as service providers are in the position to offer flexible and innovative services. The following scenario gives an example for such new services.

A telecommunication operator runs a service platform, which allows integration of 3rd party services in a flexible manner. This is used by a regional service provider (SP) to offer a novel tourist guide service (TGS) for the local city. The service can only be discovered by users who are in this city. The TGS is context-aware and thus provides enhanced user experience. For this, the TGS has to know personal attributes from its users. For example, user location and orientation is required in order to provision appropriate information about nearby famous buildings. If user's preferred language is provided then the service presents the information in this language. For some buildings videos with additional historic insights are available. To view these, the current access network QoS capabilities and user QoS parameters are considered. Because the user subscribed to the golden QoS-class and the access network has sufficient capacity, the video is presented with highest available resolution including stereo sound.

This scenario demonstrates that (1) there are short-term relations between users and service providers, (2) service providers require personal data, (3) telecommunication operators have to maintain complex business relations, and (4) users require mechanisms for attribute exchange with service providers while maintaining their privacy. To cope with these challenges the Daidalos architecture [13] is divided into two layers as shown in Figure 1.
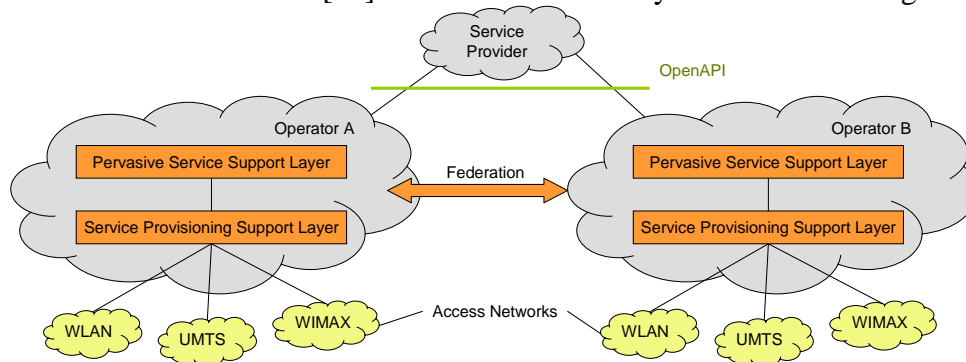

Figure 1: The Daidalos Architecture

The Service Provisioning Support (SPS) layer offers basic functionality for service deployment. This comprises, e.g. support for mobility, quality of service, multimedia services and basic security functionality, such as access network authentication. These are complemented by context and personalisation in addition to security, privacy and identity management functions in the Pervasive Service Support (PSS) layer.

4. Daidalos IdM Concepts

In Daidalos, we generalise the understanding of identity concept. It covers users, operators as well as service providers. To avoid confusion we introduced the term *entity* and limited it to natural and legal persons whose privacy has to be protected. Furthermore, a partial or virtual identity is "a subset of attributes of a complete identity" which "(each) represents the person in a specific context or role" [14]. In Daidalos the term VID is adopted. The previously referred "subset of attributes of a complete identity" thus reflects the set of personal data related to a VID and may comprise later related data items as well. Consequently, the management of VIDs becomes a key functionality to privacy protection. In the remainder, the major focus is on virtual identities with respect to users.

Our virtual identity model reflects the inherent distribution of personal data across various systems and several administrative domains as well as the association of data to virtual identities. Figure 2 shows that each entity can own several *entity profile parts* (EPPs). An *EPP* is a set of attributes that comprise a consistent whole and thus can be regarded useless if decomposed. EPPs of one entity are virtually aggregated to the *entity profile* (EP). An *entity profile* comprises all existing EPPs of an entity. Based on the *entity profile* multiple views can be defined. Each *entity profile view* (EPV) corresponds to a VID. This allows creation of several views (VID) on existing data (EPP). For more information on the Daidalos VID concept, please consult [15].


Figure 2: The Daidalos Virtual Identity Model

5. Daidalos Identity Management Architecture

In the following, the Daidalos Architecture introduced in section 3 will be detailed with respect to IdM and application of the VID model. Figure 3 shows the components of the IdM implementation, which is the basis for subsequent explanations. The components perform a two-stage process to protect user privacy and improve user convenience.

In the first phase (service preparation phase), a VID matching the characteristics and requirements of the service is selected. In the second phase (service consumption and attribute retrieval phase), the service provider can retrieve user attributes. In the following both phases are explained in more detail, including the involved components.
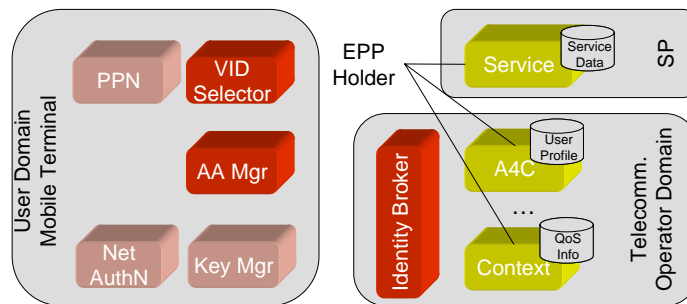

Figure 3: Daidalos IdM Components

5.1 Service Preparation Phase

The goal of this phase is authentication of the user towards the SP and to grant access to user's attributes, which are required for service provisioning.

Conventional systems describe the required user attributes in static privacy policies. This does not consider the trustworthiness of the service providers and that additional information is disclosed during retrieval of privacy policies. For example the IP address could be used to infer more details about the requesting user.

Therefore, the Privacy Policy Negotiation (PPN) component anonymously negotiates and agrees on a privacy policy [16]. A privacy policy contains statements about required attributes, disclosure to other entities as well as the purpose of processing and storage.

After successful negotiation, the VID-Selector takes the agreed privacy policy as input. It evaluates the statements about required data. Then it decides which VID of the entity in question should be used in order to fulfil the service requirements while protecting privacy. This functionality enhances the user experience in pervasive environments [17] in which a much higher dynamic with respect to changing business partners is expected.

Up to now, no personal information is revealed to the service provider. All interactions are anonymous and do not result in invasion of privacy. If the user or the VID-Selector acting on his behalf agrees to the requirements of the SP, the SP gets informed about the user's VID. This is achieved by the VID-Exchange protocol defined by the Daidalos project. This protocol binds the previously agreed privacy policy to the involved VIDs in a non-repudiable manner.

Eventually, the selected VID is activated. First, the VID is authenticated against the network. Second, the service provider must be granted access to required EPPs. This is the task of the Access Control Manager. It performs a two stage policy deployment process. In the first stage, access control rules are only deployed on the ID-Broker, which is the central hub for EPP retrieval. In the second stage, access control rules are deployed on the actual storage (EPPHolder). After successful completion of these steps the service consumption can take place.

5.2 Service Consumption and Attribute Retrieval Phase

The SP requires user attributes for the actual service provisioning. In the following, the user attribute retrieval is explained in detail.

Every VID has one ID-Broker assigned. Each ID-Broker acts as the central hub for attribute retrieval and is supposed to be operated by a telecommunication operator. Therefore, the SP queries the VID's ID-Broker to retrieve attributes (EPPs). The ID-Broker itself does not store EPPs but manages references to EPPs stored in EPPHolders (EPPH).

The ID-Broker provides two modes for attribute retrieval to support administrative domains and highly dynamic EPPs. Figure 4 and Figure 5 illustrate both modes, Proxy Mode and Refer Mode.
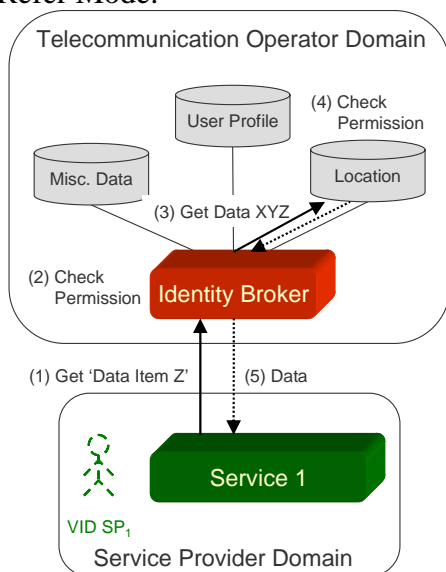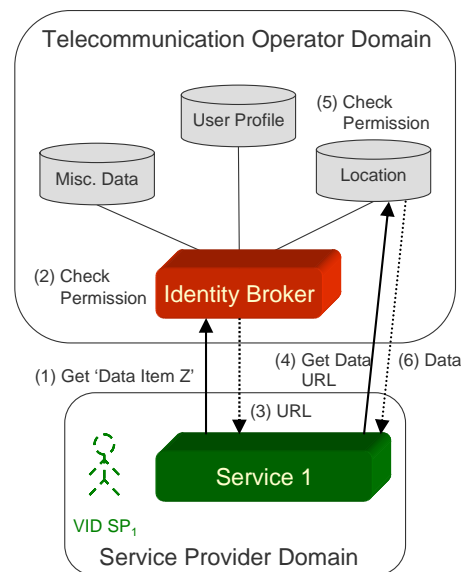

Figure 4: Proxy Mode

Figure 5: Refer Mode

Both modes have in common that the initial EPP query (1) is directed to the ID-Broker, which performs the first permission check (2).

In case of Proxy Mode, the ID-Broker contacts the EPPH (3) which optionally checks the permissions depending on the trust relationship between the ID-Broker and the EPPH. Finally, the ID-Broker receives the EPP and forwards it to the SP (5). The EPP can be provided in encrypted form to avoid that the ID-Broker can learn information about the user it is not privy to.

In case of Refer Mode, the ID-Broker only returns a reference to the actual EPPH (3). The reference allows inferring the required protocol to contact the EPPH (4). The EPPH performs an access control check (5), which is required because of possible reference caching, and returns the queried information (6).

The Refer Mode is faster if the same EPP is accessed several times, because the ID-Broker is only involved in the first access (reference caching). It also avoids the ID-Broker as a bottleneck. A drawback is that the actual storage location of the EPP is revealed. This information might be used to infer additional information about a VID and potentially link two VIDs, e.g. in case the revealed reference contains entity-specific information.

6. ID-Broker Implementation and Performance

The presented concept is prototypically implemented in the Daidalos project. In the following, details about the ID-Broker implementation including performance measurements are provided.

6.1 – ID-Broker Software Architecture

We built our ID-Broker on top of existing standards, like SAML [7] and XACML [18], to ensure a maximum compatibility with other identity management solutions.

In our solution, the ID-Broker and the EPPH are combined with required IdP functions. Our ID-Broker can thus be deployed as an EPPH enforcing access control, as an IdP, or as a "proxy" that redirects requests to the appropriate EPPH or to other ID-Brokers.

Figure 6 shows the ID-Broker internal architecture including the SAML engine and VID data model implementation (instantiated in a database).
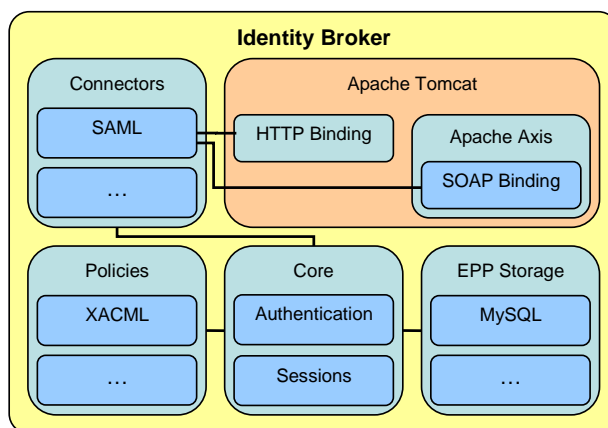

Figure 6: ID-Broker Software Architecture

Connectors allow for different protocols to communicate with the ID-Broker. We currently provide a SAML 2.0 connector that implements both SAML's HTTP and SOAP bindings, which are executed as web applications within Apache Tomcat. All requests are forwarded in a common format to the *Core* component, which keeps authentication and session state.

The Core also triggers relevant policies for each request that determine how to handle the request. We support different Policy Engines to be plugged in, and we are currently using an XACML engine. EPP resolution is delegated to the *EPP Storage* module (realises the EPPH), which contains plug-ins that translate identity information from existing components to a common format. We currently have plug-ins that provide local storage in a MySQL database and that translate data from an existing 3GPP Home Subscriber Server (HSS). Expansion of references to remote EPPs is handled by the Core, which uses a suitable Connector to expand the reference.

6.2 – Identity Broker Performance

These values were measured on a Linux-based system with kernel 2.6.22. The host machine was an AMD Opteron clocked at 2410 MHz, 512 MB RAM and 256 MB swap. We used Sun's Java SDK 1.6.0_03-b05, Apache Tomcat 5.5.23 and MySQL 5.0.45.

Table 1: Average Response Time for 100 Requests (ms)

Run	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10
AttributeQuery	173.58	156.70	173.27	151.24	170.13	163.73	152.47	146.07	139.35	149.97
AuthnRequest	146.67	179.81	167.00	147.45	139.01	133.21	166.67	160.10	156.34	140.39

Each measured time includes the sending and receiving of the request over the network, and each response produced includes an XML Digital Signature, which is where most of the time is spent. From these figures, the current implementation can process about 6 to 7 requests per second; this can be improved by distributing the load among more ID-Brokers making full use of the distributed nature of the protocol. There is also ample room for improvement in the current implementation; these metrics should be interpreted as meaning our approach is feasible.

7. Conclusions

We presented a design for a layered privacy and identity management system. It integrates pervasive services into telecommunication operator platforms. The system puts the telecommunication operator in the middle of all IdM processes. He acts as the central hub for information exchange about users without having global views on them. The user can take full advantage of all identity related information across administrative domains without entering information several times. Privacy invasion is prevented by an automated process, which selects virtual identities on behalf of users and deploys access control rules. In consequence SPs only have a restricted but sufficiently detailed view on their customers.

During system design we faced several challenges. First, we had to solve the problem of data ownership. Second, we had to deal with different representations of personal information in different layers. This is related to different understanding of terminology and different semantics of available information.

Our first implementation shows that our architecture can improve privacy and identity management in open pervasive telecommunication platforms for the operator of the platform, the user of the platforms as well as for connected service providers. Although the architecture is based on well-known IdM standards, interoperability with existing IdM solution is not automatically ensured. In future, support for groups of identities as well as improvement of the existing processes should be addressed.

Acknowledgements

References

[1] Weiser, M., The Computer for the Twenty-First Century, Scientific American, Vol. 265, pp. 94-104, 1991.

[2] Directive 95/46/EC of the European Parliament and of the Council of 24[th] October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data.

[3] Directive 2002/58/EC of the European Parliament and of the Council of 12[th] July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications).

[4] Girao, J., et al., Preserving privacy in mobile environments with virtual network stacks, GLOBECOM 2007, Washington DC, USA, November 2007.

[5] Shibboleth Architecture Technical Overview, Working Draft 02, June 2005.

[6] Liberty ID-FF Architecture Overview, Liberty Alliance Project, Version 1.2-errata-v1.0.

[7] Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0, OASIS standard, March 2005.

[8] Shibboleth Attribute Release Policy Editor, MAMS project.

[9] Liberty ID-WSF Data Services Template, Liberty Alliance Project, Version 2.1.

[10] Chappell, D., Introducing Windows CardSpace, April 2006.

[11] Policy Evaluation, Enforcement and Management Architecture (PEEM), Open Mobile Alliance standard, March 2006.

[12] 3GPP Generic User Profile (GUP); Architecture (Stage 2), Release 6, 3GPP TS 23.240 V6.7.0, March 2005.

[13] Aguiar, R. L., et al, "Pervasiveness in a competitive multi-operator environment: the Daidalos project", IEEE Communications Magazine, Large Projects section, vol 45 n.10, pp 22-26, Oct 2007.

[14] Pfitzmann, A. (editor), Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management – A Consolidated Proposal for Terminology, Version v0.30, 26[th] November, 2007.

[15] Sarma, A., et al., "Virtual Identity Framework for Telecom Infrastructures", Springer Wireless Personal Communications, Special Issue on "International Mobile Telecommunications – Advanced, 2008.

[16] Dolinar, K., et al., Pervasive systems: enhancing trust negotiation with privacy support, In Proceedings of International Workshop on Research Challenges in Security and Privacy for Mobile and Wireless Networks (WSPWN06), Miami, Florida, March 2006.

[17] Clarke, J., et al., Security and Privacy in a Pervasive World, Proceedings of the Eurescom Summit 2005, Heidelberg, 2005.

[18] eXtensible Access Control Markup Language (XACML) 2.0, OASIS standard, 2005.