Available online at www.sciencedirect.com
ScienceDirect
Computer Networks xxx (2007) xxx–xxx
Computer Networks
www.elsevier.com/locate/comnet

A middleware architecture for privacy protection

Georgios V. Lioudakis [a,*], Eleftherios A. Koutsoloukas [a], Nikolaos L. Dellas [a], Nikolaos Tselikas [a], Sofia Kapellaki [a], George N. Prezerakos [a,b], Dimitra I. Kaklamani [a], Iakovos S. Venieris [a]

[a] National Technical University of Athens, School of Electrical and Computer Engineering, 9 Heroon Polytechniou str., 15773 Athens, Greece
[b] Technological Education Institute (TEI) of Piraeus, Department of Electronic Computing Systems, 250 Thivon Avenue and Petrou Ralli, 122 44 Athens, Greece

Abstract

The issue of user privacy is constantly in spotlight since an ever increasing number of online services collects and processes personal information from users, in the context of personalized service provision. In fact, recent advances in mobile communications, location and sensing technologies and data processing are boosting the deployment of context-aware personalized services and the creation of smart environments; but at the same time, they pose a serious risk on individuals' privacy rights. Although technology makes the collection of data easy, its protection against abuse is left to data protection legislation. However, the privacy requirements, other than being general and abstract terms to be regarded as legislature issues, should be brought down in technological reality and carefully accounted for in devising technical solutions. In order to limit the disclosure and avoid the misuse of personal data, this paper discusses an architectural proposal for a middleware system that will enforce protection of user privacy through technical means. The proposed architecture mediates between the users, the service providers and the law, constituting a middleware shield for individuals' personal data.

Keywords: Privacy middleware; Privacy regulations; Policies; Ontologies; Privacy language

1. Introduction

Over the last years, the potential impact of contemporary Information and Communication Technologies on user privacy rights is regarded as being among their most evident negative effects. More than a century after the first essay identifying that privacy, as a fundamental human right, was endangered by technological advances [1], never before in history have citizens been more concerned about their personal privacy and the threats by emerging technologies [2].

This heightened awareness for privacy issues is mainly due to the ubiquity, the invisibility and the processing power of computation, communication and monitoring devices which make up the

* Corresponding author. Tel.: +30 2107722423; fax: +30 2107721092.
E-mail addresses: gelioud@icbnet.ntua.gr (G.V. Lioudakis), lefterisk@icbnet.ntua.gr (E.A. Koutsoloukas), ndellas@icbnet.ntua.gr (N.L. Dellas), ntsel@icbnet.ntua.gr (N. Tselikas), sofia@icbnet.ntua.gr (S. Kapellaki), prezerak@teipir.gr (G.N. Prezerakos), dkaklam@mail.ntua.gr (D.I. Kaklamani), venieris@cs.ntua.gr (I.S. Venieris).

1389-1286/$ - see front matter © 2007 Published by Elsevier B.V.
doi:10.1016/j.comnet.2007.06.010

distributed service provisioning infrastructures. With a densely populated world of smart and intelligent but invisible devices, no single part of the average person's life will by default be able to seclude itself from digitization. The oncoming mass deployment of context-aware and personalized services [3] and advanced monitoring and surveillance techniques [4] aiming at improving public safety demand, collect, store and process a large amount of personal data. Moreover, data mining [5] which promises to efficiently discover valuable, non-obvious information from large databases is very vulnerable to misuse and may compromise privacy by combining personal data from heterogeneous resources. Privacy related concerns are also raised due to the fact that the service provisioning chain has become increasingly complex and includes numerous actors such as service providers, content providers, operators, virtual operators, virtual service providers, service administrators and financial institutions (just to name a few) all of which can be separate entities.

So far, the vague notion of privacy is situated in the realms of legal and social studies. Nevertheless, all laws and legislation, as well as privacy codes require enforcement, the origins of which may be two-fold: self-regulation of the corresponding organizations and industry that collect and process personal data, on the one hand, and deployment of the technical means for enhancing privacy on the other. Self-regulation concerns the restriction of practices according to fair information principles. However, it is often seen as a bothersome bit of overhead, both economical and administrative, while monitoring and verification is needed in order to be effective. Besides, there are numerous cases where privacy-invasive technologies are in practice opposed to well-stated privacy policies, e.g. [6,7]. Consequently, apart from privacy protection by legislation and codes of conduct, the enforcement of legal requirements by means of privacy enhancing technologies is very important. For that reason, the European Commission, via its leading relevant Directive [8], encourages the development and use of privacy enhancing technological measures as an essential complement to legal means.

Very often, the protection of privacy is considered equivalent to security. The truth is that information and network security mechanisms, such as encryption or access control, constitute the bottom line for personal data protection. However, they certainly are not a panacea; for instance, security means cannot prevent one company from selling customers' profiles to another. Furthermore, sometimes there is even a conflict between security and privacy; e.g., by monitoring events that occur in a system, the privacy of its users may be threatened [9].

In this paper, a technical framework conceived on the basis of privacy legislation is presented. The main concept of the introduced architecture is the deployment of a unit of trust, which acts as a three way privacy mediator between the law, the users and the service providers. The central unit of the proposed middleware architecture is a privacy proxy, namely the Discreet Box, which constitutes the point where these three actors meet. The Discreet Box undertakes the enforcement of privacy regulations, by integrating into a privacy-proof middleware entity the privacy-critical operations that constitute part of the service provision chain, as well as the privacy regulations themselves, in a codified manner. Several additional components assigned with supporting functional tasks complement the architecture, while the formal definition of personal data types, as well as of services that consume personal information is attempted.

The rest of this paper is organized as follows: Section 2 provides some insights into the legal aspects of privacy. Section 3 presents the proposed middleware framework and architecture, while Section 4 provides use case examples of privacy protection using this architecture. Section 5 briefly examines related work carried out for privacy protection. The paper concludes in Section 6 with a few summarizing remarks and a short description of current work.

2. Legal aspects of privacy – fundamental requirements of a large-scale privacy system

Privacy is recognized as a fundamental human right by the Universal Declaration of Human Rights of the United Nations [10]. It is protected by relevant legislation in all the democratic countries throughout the world.

The first data protection act, adopted in 1970 by the West Germany state of Hesse, set in motion a trend towards adopting privacy legislation. The first influential text was the US Privacy Act [11], adopted by the Congress in 1974. Nowadays, the European Directive 95/46/EC [8] enforces a high standard of data protection and it is the most influential piece of privacy legislation worldwide, affecting many

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

countries outside Europe in enacting similar laws. The Directive, in effect since 1998, requires all member states to introduce necessary legislation in order to protect the right to privacy, with respect to the collection, processing, storage and transmission of personal data. Among the objectives of the Directive is the free flow of personal data between the European countries, as well as the restriction of personal data transfer only to countries outside Europe that have enforced an appropriate level of data protection.

The Directive reflects fundamental privacy principles, as codified by the Organization for Economic Co-operation and Development (OECD), in 1980 [12]. This codification was a significant milestone, as OECD principles lay out the basis for the protection of privacy. Moreover, technological advances pave the way to relevant legislation to adopt new arrangements, for conforming to the new reality. In Europe, the Directive 95/46/EC is particularized and complemented with reference to the electronic communication sector by the Directives 2002/58/EC [13] and 2006/24/EC [14], which impose explicit obligations on network and service providers to protect the privacy of users' communications. Furthermore, there is a number of official EU Opinions, Working Documents and Studies that refer to technological advances, such as the use of biometric features, high-tech surveillance mechanisms and RFIDs. Similarly, in the USA, the Computer Matching and Privacy Protection Act of 1988 [15] amended the Privacy Act by adding certain protection for the subjects whose records are used in *matching programs*.

With respect to the lawfulness and fairness of personal data collection and processing, the fundamental privacy principles may be summarized as follows (based on [16]). At the same time, these principles have to signal and constitute the basis for the definition of the functional requirements for a system that manages privacy on large scale:

- *Purpose specification and purpose binding*: The purposes for which the personal data are collected and processed should be specified and legitimate. The subsequent use of the personal data is limited to those specified purposes, unless there is an informed consent by the user. From the Privacy System's point of view, the system itself should provide the means for the explicit specification of the purpose, every time personal data are collected or processed. Moreover, the system should be able to control how the purpose is bound to the data and the action of their disclosure and processing, while the user should be able to determine the fate of his personal data.
- *Necessity of data collection and processing*: The collection and processing of the personal data should only be allowed, if it is necessary for the tasks falling within the responsibility of the data processing entity. Therefore, the system should be able to examine whether the collection or processing of specific personal data is necessary for the provision of the service in question. Additionally, it should provide – where applicable – mechanisms for the further limitation of data disclosure.
- *Information, notification and access rights of the users*: The users have the right to information, to notification and the right to correction, erasure or blocking of incorrect or illegally stored data. These rights should not be excluded or restricted by a legal transaction. In that respect, the system should provide the means for informing and notifying the users or ask for their consent, whenever some processing or disclosure of their personal data is about to take place. The users should be provided by the means for updating or deleting their personal data and determining and modifying their wish on their treatment.
- *Security and accuracy*: Appropriate technical and organizational security mechanisms have to be taken to guarantee the confidentiality, integrity, and availability of the personal data. From the Privacy System's perspective, the system – by itself – should be secure, in order to guarantee the confidentiality, integrity and availability of the personal data.
- *Supervision and sanctions*: An independent Privacy Authority has to be designated and should be responsible for supervising the observance of privacy provisions. In the event of violation of the provisions of privacy legislation, criminal or other penalties should be envisaged. In that respect, the Privacy System should provide to the Privacy Authority the means for supervising and controlling every action of personal data collection and processing.

Data protection legislation worldwide, where available, naturally defines some exceptions, exemptions and restrictions concerning the scope of the aforementioned principles. In the general case, for purposes of national security and defence, public

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

security, the prevention, investigation, detection and prosecution of crimes and other reasons of common advantage, the collection and processing of personal data may be enforced by the authorities. Lawful interception is currently a common denominator for all the regulatory frameworks for the protection of privacy and it constitutes an additional requirement that complements the list above. Therefore, the Privacy Authority should be provided with the means for the lawful interception of the personal data. However, the users should be protected from the abuse of this right that authorities obtain, while the necessary "hooks" for the lawful interception should no way become available to other entities.

3. Proposed architecture

Starting from the translation of regulatory requirements regarding handling of personal data into technical requirements for a privacy protecting system, this section provides our proposed approach in developing a middleware framework for personal data handling. The middleware architecture described here mediates between service providers and users and constitutes a distributed unit of trust that enforces the legal requirements, as far as the afore-described privacy principles are concerned.

Further in this section, Section 3.1 defines key concepts for the system's design, while Section 3.2 describes the design approach and gives a component-based view of the proposed middleware architecture.

3.1. Formal considerations

Prior to presenting the proposed middleware architecture, certain abstract notions must be formally expressed in order to establish a common ground of concepts and terminology. These notions are the *nature of personal data*, *collection and processing purposes*, *users' preferences on personal data treatment* and *privacy regulations*.

In order to determine how the proposed framework will handle personal data, it is quite essential to clarify their different types and characteristics. Therefore, a taxonomy of personal information is considered, according to some special, as well as fundamental attributes. We may identify three different categories of personal data, which pose certain design requirements for the proposed framework; the architecture must be capable of handling effectively all three types of personal data:

- Certain personal data may be characterized as "active", in the sense that the user actively controls them. It is up to the user whether to disclose his identity or his credit card number when prompted. An essential characteristic of these data is that their source is the user's terminal.
- Personal data that originate from sensors and RFIDs can be characterized as "semi-active". As far as such data are concerned, the user has only partial control over them. These data are not actively provided by the user, but rather extracted from tokens. However, in the general case, the semi-active data are directly linked to the user's identity.
- Personal data that are produced and disclosed without any specific user's action are termed "passive". Surveillance video constitutes a typical example. In this case, the user only passively consents to the collection of these data, while in most cases, he is not even notified.

Apart from the category where a piece of personal information falls under the taxonomy described above, its specific type constitutes a critical parameter for its treatment by a privacy enhancing framework. In [12], this is defined as "different degrees of sensitivity". Medical records and credit card numbers are considered more sensitive than some abstract service preferences, as far as privacy issues are concerned. Consequently, a semantically uniform formal way of describing the specific type of every possible piece of personal information must be considered. Therefore, in this paper, the *Ontology of Personal Information* is defined. It constitutes an ontology [17] that serves as the common information model that is shared across the system and enables its modules to interoperate. This way, when performing any kind of processing on some personal data, the processing module, either privacy enhancing or service provision related, is capable of precisely understanding the type of the information and act accordingly based on the predefined rules which have been set for that particular type.

Since the critical issue tackled by the middleware is personal data handling, focus needs to be put on the different uses of the structured personal information types. Therefore, a classification of the potential services, similar to what the Common Procurement Vocabulary [18] represents for public pro-

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

curement in Europe, must be defined. In that respect, a second ontology is defined, the *Ontology of Services*, which formally describes the services for the provision of which personal data are demanded, collected, disclosed and processed. Fig. 1 illustrates indicative segments of the two ontologies.

Finally, what remains is to formalize how users can express their privacy preferences and how regulations can be expressed in the system. The proposed approach is to use a policy framework that will incorporate a large set of rules. Each regulatory requirement might be implemented in a series of rules that permit/forbid operations based on the enabled privacy protection measures for every specific type of personal information and the declared intended use. Default rules inherent in the system will provide for the regulatory requirements, while user privacy preferences can be incorporated at the policy framework and be considered during the decision making process. Rules that originate from users provide some flexibility and the sense of control of privacy to users.

In order to express the rules, either user-or regulations-originated, in a formal way, we define a relevant XML-based [19] language, namely the Discreet Privacy Language (DPL). The formal definition of the DPL is beyond the scope of this paper; however, the description of some basic characteristics is provided in the following.

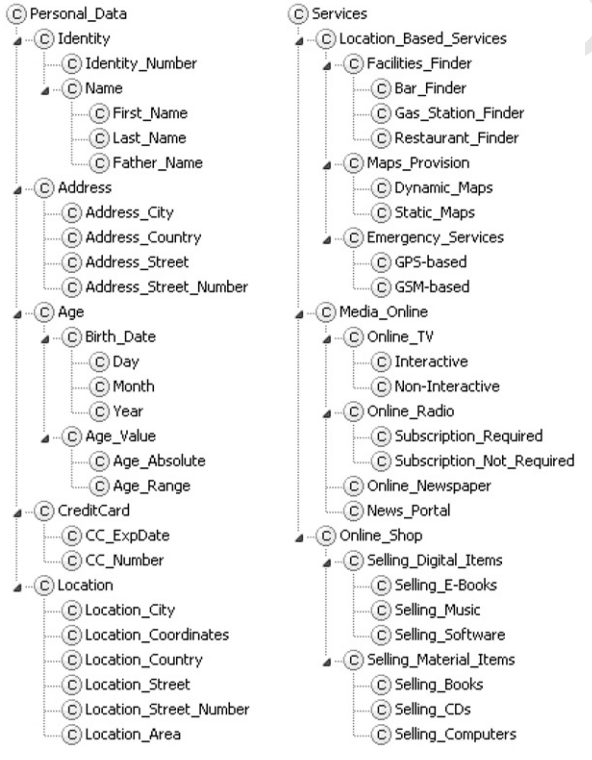

Fig. 1. Ontology segments.

The user-and regulations- originated rules' statements are enclosed into `PREF_RULE` and `REG_RULE` elements, respectively. Several sub-elements may be included inside a rule statement:

- `DATA_TYPE`: Expresses the type of the data in question; its values come from the domain defined by the Ontology of Personal Information. It can also take the value `ALL`, covering all the types of data.
- `DATA_TYPE_DESCENDANTS`: Denotes whether the defined rule is applicable by inheritance to the descendants of the specified `DATA_TYPE` in the class hierarchy of the Ontology of Personal Information. It can take the value `YES` or `NO`.
- `SERVICE_TYPE`: Expresses the type of the service for which some data are about to be disclosed or processed; its values come from the domain defined by the Ontology of Services. It can also take the value `ALL`, covering all the services.
- `SERVICE_TYPE_DESCENDANTS`: Denotes whether the defined rule is applicable by inheritance to the descendants of the specified `SERVICE_TYPE` in the class hierarchy of the Ontology of Services. It can take the value `YES` or `NO`.
- `RULE_TYPE`: Denotes the subject of the specified rule. Some typical values for the `RULE_TYPE` are:
  - `DISCLOSURE_Y-N`: It determines whether the data of `DATA_TYPE` for the service of `SERVICE_TYPE` should be disclosed to the provider or not.
  - `DISCLOSURE_LEVEL`: It determines the level of abstraction for the data of `DATA_TYPE` for the service of `SERVICE_TYPE`.
  - `NECESSARY_Y-N`: It determines whether the data of `DATA_TYPE` is necessary for the provision of a `SERVICE_TYPE` service.
  - `MODIFICATION_PERMISSION`: It determines whether the service provider has modification privileges of data of `DATA_TYPE` during the provision of a `SERVICE_TYPE` service.
  - `NOTIFICATION`: It determines whether the user should be notified for some action on his data of `DATA_TYPE` for the service of `SERVICE_TYPE`.

- CONSENT: It determines whether the user should be asked for his consent for some action on his data of DATA_TYPE for the service of SERVICE_TYPE.
- RETENTION_PERIOD: It specifies the retention period for the data of type DATA_TYPE.
- RESTRICTION: It sets restrictions of subject SUBJECT to the rule(s) specified inside a statement.

• VALUE: It is the value of the subject of the specified RULE_TYPE.

Section 4 provides several examples of the DPL, in the context of the use cases described there, while the detailed description of DPL along with its ABNF [20] normative specification can be found in [21].

3.2. Design approach and component view

The requirements and formalism considerations described above cater for the following design approach for privacy protecting middleware. On one hand, the Ontology of Personal Information enumerates what the system is able to handle and sets the common ground for the system's information model. On the other hand, the Ontology of Services enumerates different occasions where personal data are involved. The binding glue between the two is a set of rules that originate either from regulatory requirements or from user privacy preferences, which map services onto personal information types. This mapping is used by the system to define whether a personal data request from a service is permitted or not, according to the current privacy status, the currently enabled privacy protecting measures and the privacy level required for each particular data type. This mapping serves as a system wide pointer to (data type, service type) couple, for which a number of regulations or user preferences need to apply. This process is outlined in Fig. 2.

In order to define the components of the architecture, another key concept needs to be taken into account. If the service provider possesses some data, there is no technically feasible way for abuse prevention. This fact leads to an approach that is based on the principle of limiting or even avoiding the disclosure of personal information to the service providers. However, this basis is cancelled as soon as the data in question are stored in service provider databases; in this case, the service provider has the complete control over the data.

Solutions like storing the personal data inside the middleware or permanently at the user's terminal have proved to be non practical. Neither a mediating entity nor a terminal can afford the processing load of global service provisioning, while both approaches inevitably demand the use of mobile code [22], in order to move the service logic to the execution points. This would introduce significant security problems to the architecture [23]. Besides,

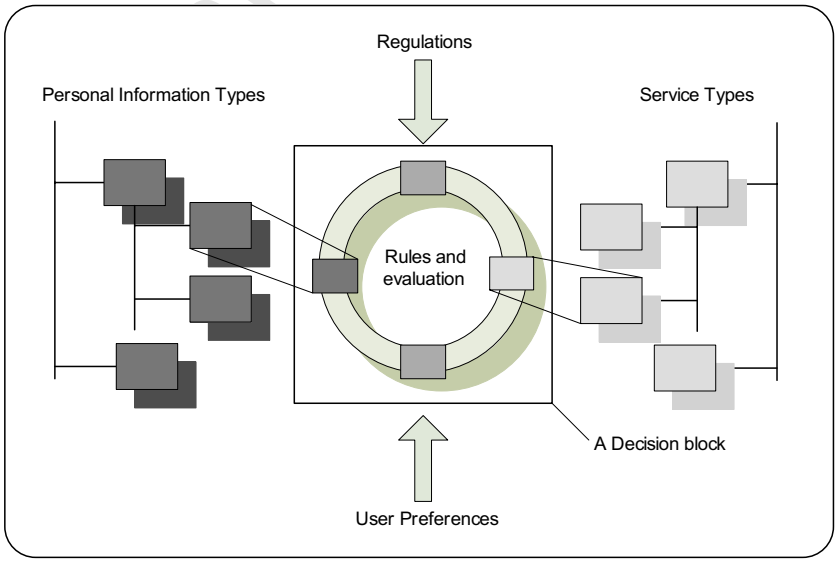

Fig. 2. Matching between personal data and service types.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

not all data are generated at the user's terminal. In fact, there are data types that are natively generated at the service provider's side, e.g., all data defined as passive.

Therefore, the proposed architecture introduces the concept of the *Discreet Box*, which constitutes a privacy proxy installed at the service provider's premises but totally controlled by the Privacy Authority. The Discreet Box incorporates the personal data repository that caches personal data and enforces retention periods, a policy framework that takes the decisions for personal data disclosure and interfaces to the service provider for the request/response cycles. It is built according to security standards that ensure the confidentiality, integrity, and availability of the personal data stored and processed inside, as well as the availability and reliability of the system itself. An analysis of the mechanisms that ensure the security and availability of this system are out of the scope of this paper, since our focus is on privacy elevation through the system's design.

The other architectural components of the proposed framework are the Privacy Infrastructure Components that form the Privacy Network for the purposes of Discreet Boxes' online monitoring, management and Lawful Interception, the User Privacy Manager which is the user-side component and the service providers' applications. Fig. 3 outlines these entities and their place in the framework.

The Discreet Box (Fig. 4) serves as the entry point to a service and its operation is similar to a proxy server. Casual interaction between the user and the service that does not involve any personal data passes transparently through the Discreet Box to the service. When the service logic reaches the point where personal data need to be involved, the service submits requests for specific data types to the Discreet Box, through a dedicated API, designating a specific purpose. The Discreet Box then performs the appropriate actions, depending on the category of data (active, semi-active, passive), in order to collect and disclose the requested data

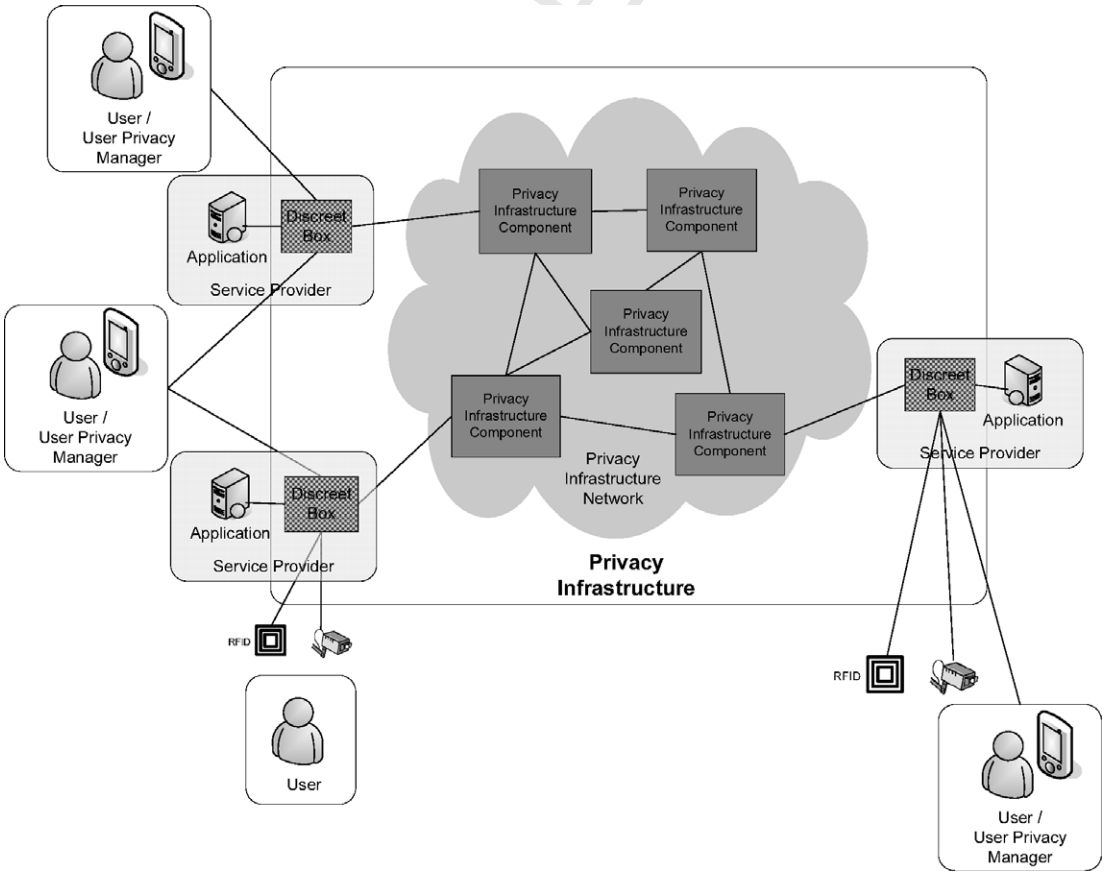

Fig. 3. Framework design.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

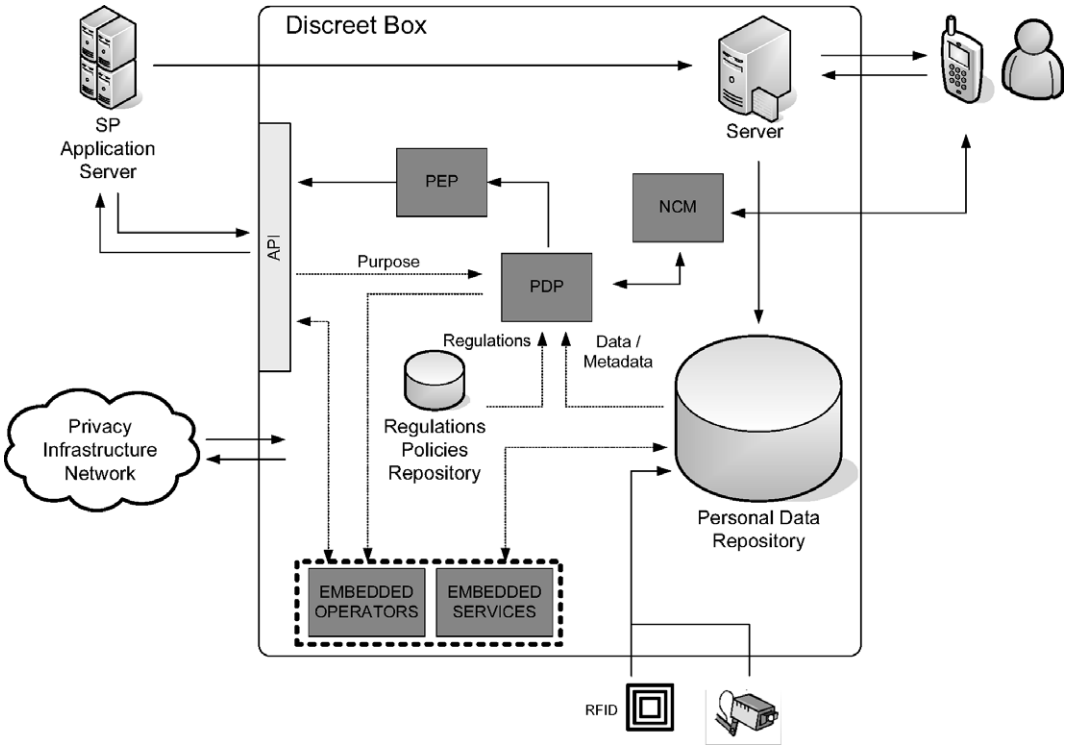

Fig. 4. Discreet Box.

permitted by the mapping between data type and purpose.

The mapping involves a number of rules originating from the regulations or the user privacy preferences; the former are fetched from the Regulations Policies Repository, while the latter are communicated by the user's terminal as metadata accompanying the respective personal data. This intelligent mapping is performed by the Discreet Box's PEP/PDP [24] policy framework that guards access to the data and interfaces to the different sources of personal information or interacting entities. The Regulations Policies Repository is always kept up to date by the Privacy Infrastructure Component to which the Discreet Box is registered.

The personal data are stored encrypted inside the Discreet Box, in the Personal Data Repository. Every action performed on the contents of the Personal Data Repository by any of the involved entities are subject to detailed logging for either real time (by associated triggers) monitoring and control of the related activities or their re-examination. The storage may be short-time (e.g., immediate service provision) or long-time (e.g., e-government services that require information archives). When the personal data in question are active or semi-active, they are stored along with their metadata, that is, the data type and user-defined privacy preferences' statements. In the case of semi-active data, the preferences are dynamically set by the user, while in the active data case, the metadata are transmitted along with the corresponding data.

In that respect, a data structure is defined, namely the Discreet Data Lock. It constitutes a shell that encapsulates the personal data transmitted by the terminal to the Discreet Box, along with their metadata in an encrypted and signed object that ensures safe communication. The Discreet Data Lock structure adds some elements to the DPL's description of Section 3.1: DATA, which encloses the lock and META, which surrounds the metadata.

The Discreet Box provides the user with the appropriate access rights in order to be able to update or delete his personal data and the associated metadata. To that respect, the Discreet Box exposes the corresponding interfaces, which are accessed by the user via the User Privacy Manager, either directly or through the Privacy Infrastructure Network.

Whenever the user's notification or consent is required for any action on personal data, the Noti-

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

fication and Consent Manager (NCM) module of the Discreet Box undertakes the task of appropriately interacting with the User Privacy Manager. It is noted that the notification may be disabled or the consent may be a priori given by the user. These are among the issues expressed in the metadata and examined by the policy framework.

The Discreet Box's internal structure is complemented by the Embedded Operations and Embedded Services modules. These modules undertake the execution of simple data processing tasks and whole services' parts respectively, in order to further reduce the amount of disclosed data. Typical Embedded Operators functionalities are the filtering of the data precision prior to their disclosure (e.g., the translation of exact location to more abstract terms) and the filtering of passive data (e.g., face blurring in surveillance video). Embedded Services concern the execution of standard service components internally (e.g., e-mail sending or service charging mediation).

The User Privacy Manager (Fig. 5) constitutes the user-side component of the architecture and resides in user's terminals. The User Privacy Manager has the two-fold purpose of being the user's assistant regarding privacy issues (e.g., it is used to set privacy preferences, edit a privacy profile, notify the user of privacy related alerts generated from Discreet Boxes, etc.) and at the same time it employs a policy framework of its own so that personal data that originate from the user (active data) are guarded even before they leave the terminal equipment. It incorporates a Notification and Consent Manager (NCM) module for the interaction with the peer module of the Discreet Box, as well as a Discreet Data Lock Constructor (DDLC), in order to create the secure data structures. The interaction with the user is performed through the Privacy Manager Console GUI.

It should be noted here that, regarding the communication facilities required to support the interaction between the entities of the privacy environment and especially the transmission of Discreet Data Locks, a SOAP based RPC mechanism is being developed. A collection of messages realizes a proprietary protocol for the interaction between components of the three central architectural entities. In that respect however, an emerging challenging issue is how to integrate the variety of user interfaces of services with a privacy console through which the user can safely provide personal data. Our approach is first to develop bindings for the HTTP protocol that is used to transfer the HTML-based service interface, and secondly to semantically mark input forms at the HTML interface with pointers to the actual personal data being transferred in a parallel SOAP-based communication stream. For services that require applications at the terminal device and do not rely on a Web interface, the Privacy Manager Console will provide a homogeneous GUI for personal data input and the required handles to the application. The detailed specification of the communication and user interface mechanism is work under way, however the reader may refer to [21] for some preliminary design results.

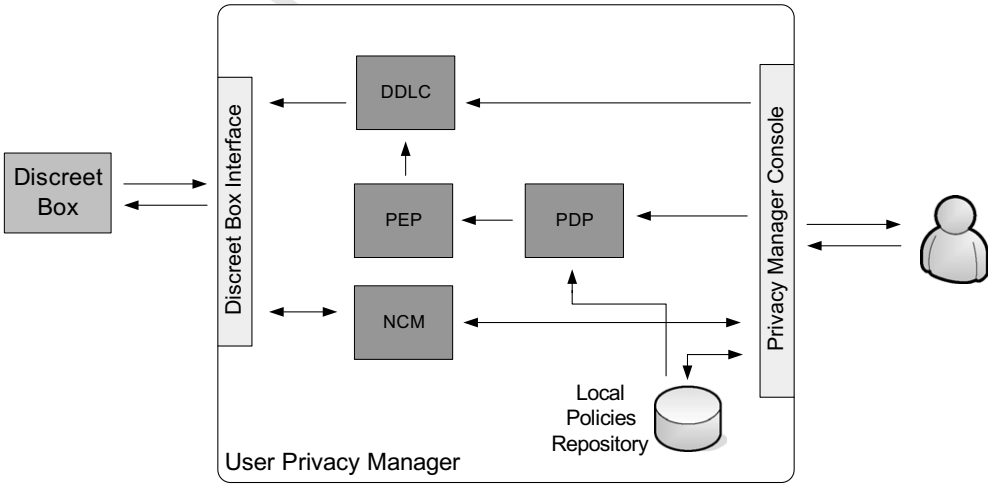

Fig. 5. User Privacy Manager.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

The Privacy Infrastructure Component, for which the logical structure is illustrated in Fig. 6, constitutes the Privacy Authority's entry point to the framework. The component supports three types of operations: Policy Authoring, Discreet Box management (e.g., regulatory policies installation) and Lawful Interception, through the appropriately defined interfaces. The component contains a policy framework for the control of the Lawful Interception and any other operation that demands access control, as well as a policy repository for the authored regulatory policies. All the Discreet Boxes registered to a Privacy Infrastructure Component are kept in the Discreet Boxes Registrar. Naturally, the Privacy Infrastructure Component provides interfaces for the communication with the Discreet Boxes and the other Privacy Infrastructure Components in the Privacy Infrastructure Network.

The Privacy Infrastructure Network constitutes a logical network of trust, through which all the privacy related functions that require interaction between the architectural entities are performed. When a regulations' modification occurs, the updated policies are disseminated to all the Discreet Boxes through this network. The tasks of Discreet Boxes online monitoring, mass user's preferences modification and Lawful Interception are among the operations where the Privacy Infrastructure Network participates. Any related activity is extensively logged by the Infrastructure Components.

4. Use case examples

This Section provides use case examples that demonstrate the proposed framework's operation during the provision of typical services.

The first use case concerns the provision of a context-aware service. We consider a service provider that offers a number of location-based services to his customers. With respect to the ontologies' segments provided in Section 3.2, the user of the use

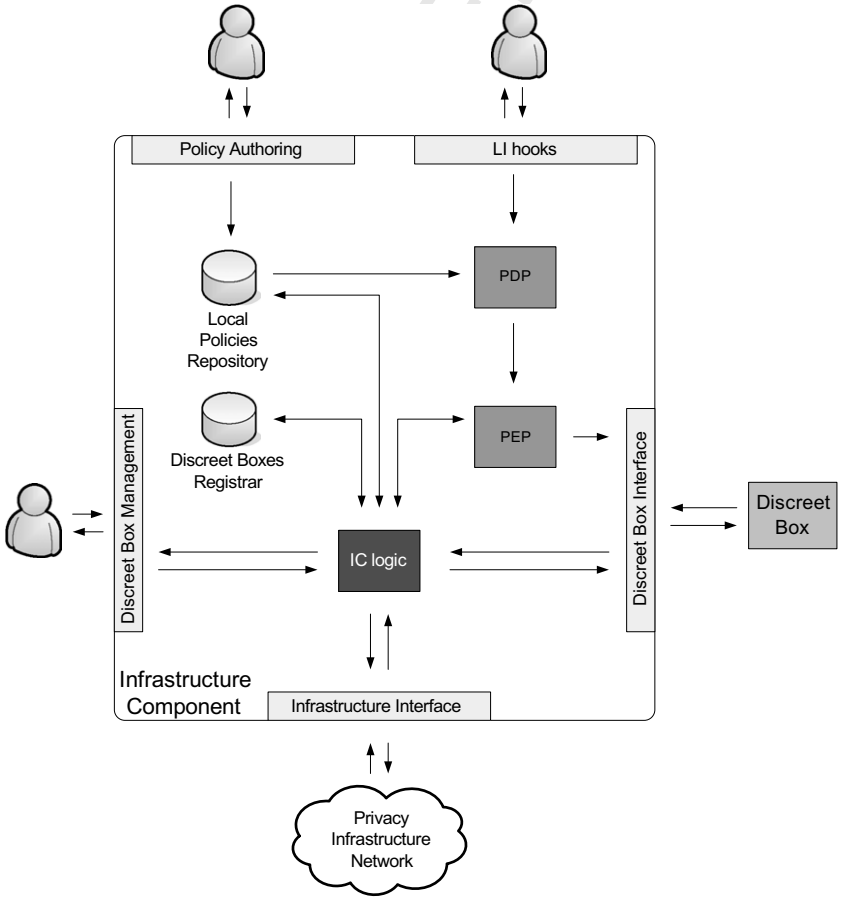

Fig. 6. Privacy Infrastructure Component.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

case is asking for the provision of the `Bar_Finder` service: he asks for the jazz bars around his location. The personal data needed for the service provision are the user's location, identity and credit card information for the service's charging and his age, supposing that information related to bars is provided only to adults. We assume that the user is a returning, registered customer; that is, his identity, credit card information and age are already stored in the Personal Data Repository of the provider's Discreet Box, while his location has to be provided in real time. Fig. 7 illustrates the service provision steps.

The procedure begins with the user requesting the service. We suppose that the user provides at this time his credentials, in order to be identified as a registered customer. The request is relayed to the service provider, without disclosing the identity of the user. Upon receiving the request, the service provider asks for the needed data; this is done through the API offered by the Discreet Box.

The Discreet Box checks for the validity of the data request, i.e., the necessity of data collection and processing and the specification and binding of the purpose. In that respect, the rules governing the personal data and service in question are retrieved from the Regulations Repository, in order to be examined. The corresponding rules indicate that since the `Bar_Finder` service is a descendant of the service class `Location_Based_Services`, the user's location is indeed needed, while the age is needed as well, for the adult age verification demand of the requested service (Fig. 8a). Therefore, the user is asked for the former, while the latter is fetched from the Personal Data Repository. The user's position is determined by his GPS-enabled terminal, in terms of geographical coordinates. The location is transmitted to the Discreet Box after being encapsulated into a Discreet Data Lock, along with metadata expressing the predefined user's preference of using the highest possible level of abstraction when disclosing location information to service providers (Fig. 8b).

The metadata of the user's age (Fig. 8c), include his demand for his explicit consent every time when any information concerning his age is about to be

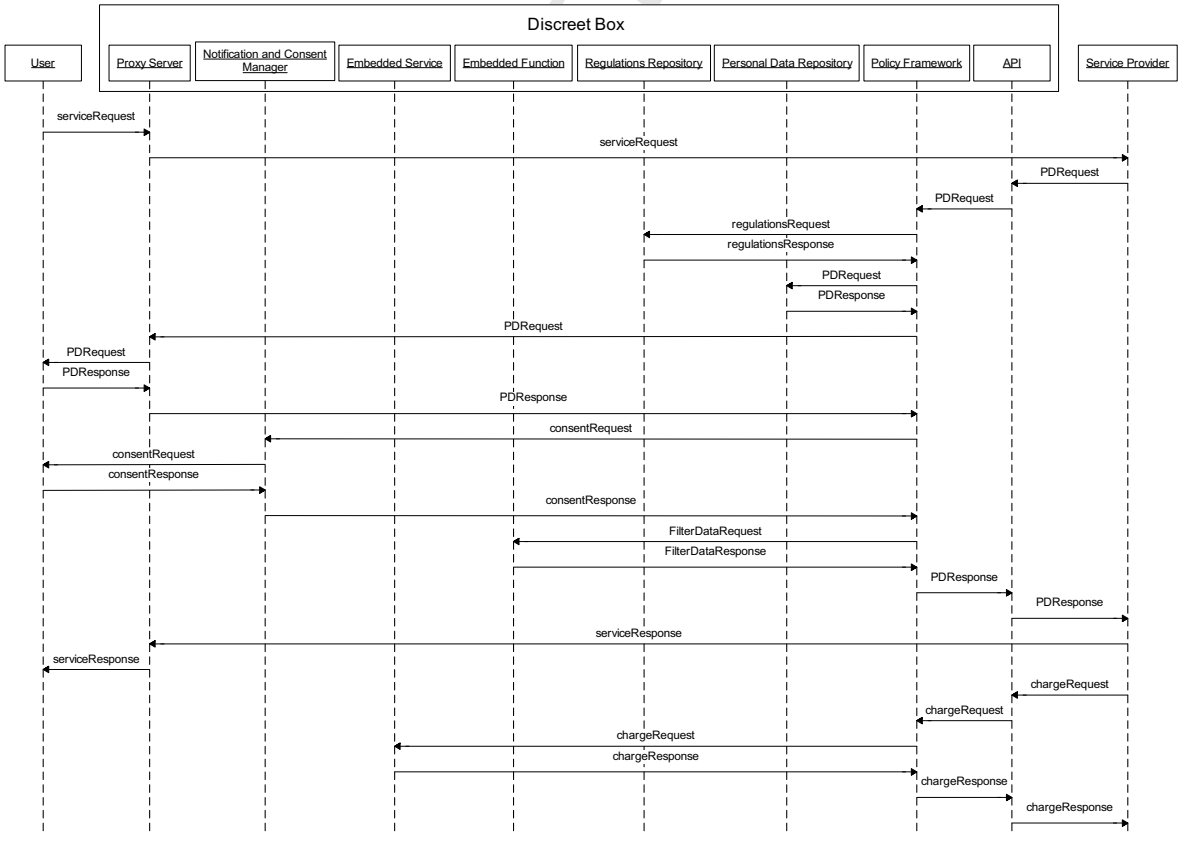

Fig. 7. `Bar_Finder` service sequence diagram.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

```
<REG_RULE>
        <SERVICE_TYPE>Location_Based_Services</SERVICE_TYPE>
        <SERVICE_DESCENDANTS>YES</SERVICE_DESCENDANTS>
        <DATA_TYPE>Location</DATA_TYPE>
        <DATA_TYPE_DESCENDANTS>YES</DATA_TYPE_DESCENDANTS>
        <RULE_TYPE>DISCLOSURE_Y-N</RULE_TYPE>
        <VALUE>YES</VALUE>
</REG_RULE>

<REG_RULE>
        <SERVICE_TYPE>Bar_Finder</SERVICE_TYPE>
        <SERVICE_DESCENDANTS>YES</SERVICE_DESCENDANTS>
        <DATA_TYPE>Age</DATA_TYPE>
        <RULE_TYPE>NECESSARY_Y-N</RULE_TYPE>
        <VALUE>YES</VALUE>
        <RESTRICTION>
                <SUBJECT>DISCLOSURE_LEVEL</SUBJECT>
                <VALUE>MINIMUM_POSSIBLE</VALUE>
        </RESTRICTION>
</REG_RULE>
```

a) Regulations' rules

```
<DATA>
        <DATA_TYPE>Location_Coordinates</DATA_TYPE>
        <VALUE>123.456789, 987.654321</VALUE >
        <META>
                <PREF_RULE>
                        <RULE_TYPE>DISCLOSURE_LEVEL</RULE_TYPE>
                        <SERVICE_TYPE>Bar_Finder</SERVICE_TYPE>
                        <VALUE>MINIMUM_POSSIBLE</VALUE>
                </PREF_RULE>
        </META>
</DATA>
```

b) Discreet Data Lock

```
<PREF_RULE>
        <DATA_TYPE>Age</DATA_TYPE>
        <DATA_TYPE_DESCENDANTS>YES</DATA_TYPE_DESCENDANTS>
        <SERVICE_TYPE>ALL</SERVICE_TYPE>
        <RULE_TYPE>CONSENT</RULE_TYPE>
        <VALUE>ALWAYS</VALUE>
</PREF_RULE>
```

c) User consent demand for disclosure of Age

```
<PREF_RULE>
        <DATA_TYPE>Age</DATA_TYPE>
        <DATA_TYPE_DESCENDANTS>YES</DATA_TYPE_DESCENDANTS>
        <SERVICE_TYPE>ALL</SERVICE_TYPE>
        <RULE_TYPE>DISCLOSURE_LEVEL</RULE_TYPE>
        <VALUE>MINIMUM_POSSIBLE</VALUE>
</PREF_RULE>
```

d) Use_r preference for disclosure level of Age

Fig. 8. Bar_Finder DPL statements.

disclosed. Therefore, the Discreet Box, via the Notification and Consent Manager entity, notifies the user and asks for the relevant consent.

At this time, the Discreet Box possesses all the data that are needed for the service provision. However, the user has expressed his preference on providing the service provider with the least possible amount of information (Fig. 8b and d). Therefore, instead of the accurate Age_Absolute and Location_Coordinates data, only the useful part of the information is disclosed, that is Age_Range to denote that the user's age is above the necessary threshold and Location_Area, pointing the broad area where the user is located, respectively.

It has to be noted here that at first sight, it might seem excessive to hide exact information of certain types (e.g., for age) in the cases where the identity is not disclosed. However, sometimes someone's identity may be inferred by such data, especially under cross-organization combinations. In [25], it is shown that 53% of the population of the USA can be uniquely identified by only {place, gender, date of birth}, where place is basically the city, town, or municipality in which a person resides. Therefore, the system strictly applies the defined policies for all data types.

Based on the given data, the service provider generates the list of the bars residing within the specified area and communicates it to the user, with the mediation of the Discreet Box. The procedure concludes with the charge of the provided service. It is assumed here that the user has provided his consent about being charged for the provision of the service, by explicitly requesting the service. The charge is performed by the associated Embedded Service of the Discreet Box. This way, the process is fulfilled without the disclosure of the user's identity or credit card information to the service provider.

The second use case concerns the automatic vehicle identification that relies on RFIDs and is deployed by many toll collection systems (e.g., New York's E-ZPASS [26]). With respect to the Ontology of Services as defined in Section 3.2, we consider two different services. The first service is simply the RFID_Electronic_Toll_Collection; its path in the ontology is Transport_Service.Toll_Service.Electronic_Toll_Collection. RFID_Electronic_Toll_Collection. It is the pre-paid service which automatically identifies the vehicle by its RFID tag, checks for remaining credits and automatically opens the toll gates. The second service concerns the notification of the user via e-mail, when his credit is below a certain threshold. The service is identified as Low_Credit_E-Mail_Notification and situated in the ontology at Pre-Paid_Service.Low_Credit_Notification. Low_Credit_E- Mail_Notification.

What essentially differentiates this use case from the context-aware one is the use of semi-active data and that the role of the user is passive; in fact, the provision of both services is fired by the toll system's application which performs a database update as well, while the user provides in a passive way the identity of the RFID tag.

The RFID_Electronic_Toll_Collection (Fig. 9) service starts with the RFID tag's response to the RFID reader's signal broadcasting [27], denoted at Fig. 9 as *RFIDResponce* and *RFIDRequest* messages, respectively. After the identification

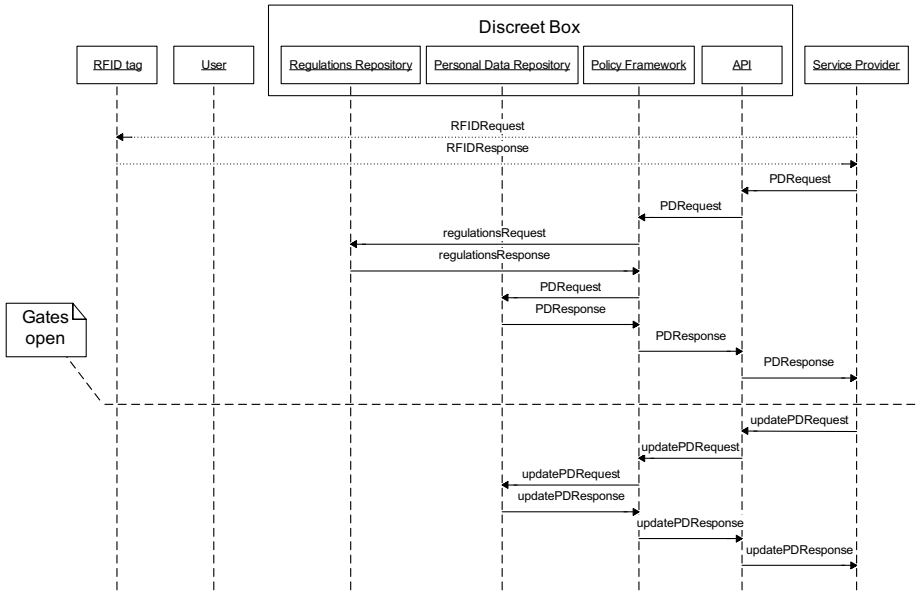

Fig. 9. RFID_Electronic_Toll_Collection service sequence diagram.

of the RFID tag, the service provider's application needs to check for the availability of remaining credits of the vehicle associated with the specific tag. Therefore, it needs the corresponding piece of information, i.e., the data of type Remaining_Credits, which is asked by the Discreet Box. The Discreet Box checks for the validity of the data request, using the same mechanism illustrated at the context-aware use case example, and infers that the data should be provided, since the Remaining_Credits data type is associated with the Electronic_Toll_Collection service (Fig. 11a), which is an ancestor of the RFID_Electronic_Toll_Collection service. Supposing that there are remaining credits, the gates open. At this time, the application wants to update the database with the transaction's detailed record and the new remaining credits. The Discreet Box permits

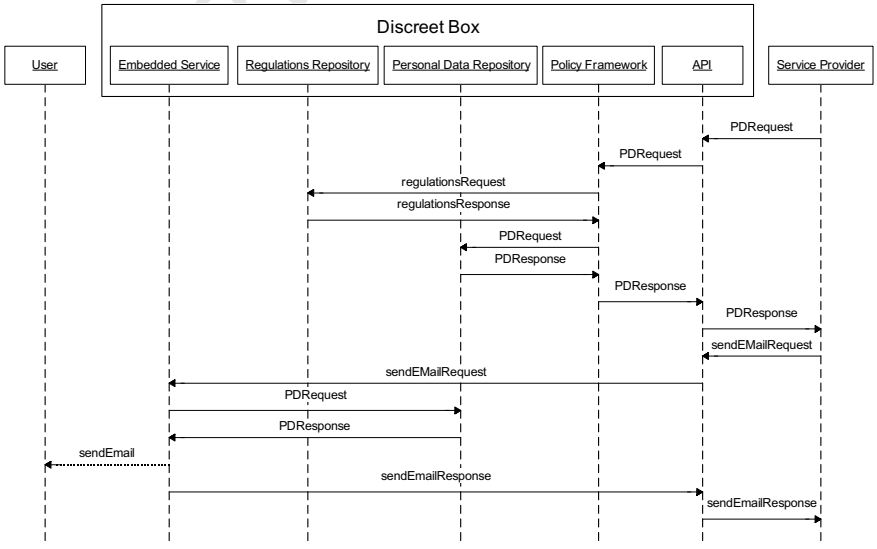

Fig. 10. Low_Credit_E-Mail_Notification service sequence diagram.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

the database modification based on the corresponding rule for the RFID_Electronic_Toll_Collection service (Fig. 11b).

Assuming that the credits of the user have reached the borderline, the service provider's application wants to notify the user via e-mail to add credit to its toll account. Therefore, it activates the Low_Credit_E-Mail_Notification service (Fig. 10). According to the associated rules (Fig. 11c), the Discreet Box will not disclose the user's e-mail address. Instead, the Embedded Service charged with the (standard) task of sending e-mails undertakes the process. The service provider's application provides the Discreet Box via the API with the RFID tag identifier, as well as with the body of the e-mail and the other relevant parameters (e.g., subject). In order to make the e-mail more personalized (e.g., to start with a phrase like "Dear Mr. John Smith") personal data are needed. This functionality is covered the use of special tags inside the e-mail which are parsed by the Embedded Service and replaced by the corresponding personal data values. To that respect, the Embedded Service fetches these data by the Personal Data Repository, along with the e-mail address of the user carrying the specified RFID tag. It is noted that this Embedded Service – Personal Data Repository transaction does not involve any interaction with the policy framework, since the e-mail and the personal data contained inside are destined to the user himself.

```
<REG_RULE>
        <SERVICE_TYPE>Electronic_Toll_Collection</SERVICE_TYPE>
        <SERVICE_DESCENDANTS>YES</SERVICE_DESCENDANTS>
        <DATA_TYPE>Remaining_Credits</DATA_TYPE>
        <RULE_TYPE>DISCLOSURE_Y-N</RULE_TYPE>
        <VALUE>YES</VALUE>
</REG_RULE>
```

a) Regulation rule for Remaining_Credits disclosure

```
<REG_RULE>
        <SERVICE_TYPE>Electronic_Toll_Collection</SERVICE_TYPE>
        <SERVICE_DESCENDANTS>YES</SERVICE_DESCENDANTS>
        <DATA_TYPE>Remaining_Credits</DATA_TYPE>
        <RULE_TYPE>MODIFICATION_PERMISSION</RULE_TYPE>
        <VALUE>YES</VALUE>
</REG_RULE>
```

b) Database modification Regulation rule

```
<REG_RULE>
        <SERVICE_TYPE>Low_Credit_E-Mail_Notification</SERVICE_TYPE>
        <DATA_TYPE>e-mail_address</DATA_TYPE>
        <RULE_TYPE>DISCLOSURE_Y-N</RULE_TYPE>
        <VALUE>NO</VALUE>
        <RESTRICTION>
                <SUBJECT>USE_EMBEDDED_SERVICE</SUBJECT>
                <VALUE>E-MAIL_SERVICE</VALUE>
        </RESTRICTION>
</REG_RULE>
```

c) Regulation rule for using Embedded Service

Fig. 11. RFID_Electronic_Toll_Collection and Low_Credit_E-Mail_Notification DPL Statements.

5. Related Work

In the last few years many approaches have been presented in order to address the issue privacy protection of individuals. The proposed solutions try to provide privacy protection by introducing identity management and/or access control mechanisms based on privacy policies, mainly focusing on enterprise environments.

Privacy policies are currently considered among the most prominent privacy-enhancing technologies. Privacy policies concern the formal specification of an organization's business practices regarding the collection and the consequent use of personal data. The privacy policies are supposed to be restricted according to fair information principles and to comply with the relevant legal framework. Privacy legislation dictates how personal data should be treated after their provision by the data subjects to service providers and other processing entities, defining in essence the requirements for the privacy-aware management of personal data through their whole lifecycle.

The Platform for Privacy Preferences (P3P) W3C specification [28] is the first initiative towards this direction, providing a way for a Web site to encode its relevant practices and to communicate them to the users that visit the site. Since it was proposed, P3P has received broad attention from both industry and research community and many complementary frameworks have been proposed (e.g. [29,30]), but it has also been subject of criticism from the current technical work, e.g. [31]. The major issue with P3P is the lack of mechanisms for the enforcement of the specified privacy policies. In essence, P3P formalizes privacy promises given to the users for fair information practices; nevertheless, after their disclosure to a service provider, there is no guarantee about the fate of a user's personal data. Besides, there are numerous cases where real practices contradict well-stated privacy policies, e.g. [6,7].

The challenge of enforcing a privacy policy has been thoroughly examined and several different solutions have been proposed, e.g., by IBM [32–37], OASIS [38] and Hewlett Packard [39–42]. These frameworks mainly focus on enterprise environments and provide the means for the automation of the privacy policies enforcement. The means for achieving this is to apply privacy-aware access control mechanisms which enhance traditional Role-Based Access Control (RBAC) [43] models with additional, privacy-related aspects, such as the pur-

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

pose for data collection, retention periods, user consent, etc.

The introduction of identity management with respect to privacy is the subject of several works, like the European Union funded projects PRIME [44] and FIDIS [45] and the Liberty Alliance project [46] which constitutes an effort by a consortium of major technology vendors and consumer-facing enterprises, formed to develop an open standard for federated network identity. The Liberty Alliance project is based on the concept of enabling users to connect multiple sets of personal information which exist across several service providers into a single easy-to-manage federated identity. While one of the key objectives of the Liberty Alliance is to enable consumers to protect the privacy of their identity information, the multidiscipline specifications Liberty covers makes it vulnerable to a variety of privacy breaches, as has been pointed out in [47,48].

All these solutions have their weak points. First, although they manage to address the issue of privacy policies internal enforcement within an organization to a great extent or to enable the privacy respectful management of user profiles and identities, they fail in providing the necessary guarantees for fair information practices to the users. In fact, since an organization possesses some personal data, their use or abuse by means of processing and disclosure are still based on good intents. Misuse may occur by a malicious employee with legitimate access to sensitive data or by any form of direct access to the data that bypasses the privacy protecting system. That is, in all these cases a user must trust a service provider. Second, the privacy policies specified in the context of these frameworks cannot be efficiently audited and verified as far as their regulatory compliance and consistency is concerned. Even an organization with the best intentions may specify a privacy policy that is not legislation-proof. Third, the specification of complex privacy policies and the continuous process of keeping them up-to-date introduce significant economical, operational and administrative overhead to an organization. In that respect, these have been the motivating principles for the framework presented in this paper.

6. Conclusion

The recent technological advances in mobile networking, sensor networks, ubiquitous and context-aware computing are reshaping the lives of individuals, facilitating it and improving its quality. However, they put user privacy at a serious risk, since they realize Ron Rivest's "reversal of defaults": what was once private is now public; what once was hard to copy is now trivial to duplicate; what was once easily forgotten is now stored forever.

In this paper, a framework for the protection of personal data is provided and studied. The central idea behind the framework is the integration of all privacy-critical functionality into a privacy-proof middleware proxy entity, namely the Discreet Box, assigned with the mission of enforcing the privacy legislation principles, on the basis on which it is conceived. The use of a privacy proxy has been exploited as a straight-forward solution, since the assignment of the complete control of the personal data to the user terminal has been proved impractical due to the limited processing capabilities of mobile devices, certain security issues and the fact that not all data are produced at the terminal. The framework, deployed by a Privacy Authority, is complemented by several assistant architectural entities, as well as by the formal definition of personal data types and services.

The benefits of using this privacy proxy can be summarized as follows. In order to provide guarantees for the enforcement of the privacy principles, the proposed middleware architecture maintains full control on the lifecycle of personal data. With the explicit separation of the personal data from the applications of the service providers, the latter cannot gain access to some personal data other than the one specified by legislation and user preferences. That is, every piece of personal information on its way to the service providers is filtered by the Discreet Box, so that the danger of unauthorized access and misuse is eliminated. The incorporation of several privacy-critical processing functionalities in the Discreet Box by means of Embedded Services and Operators further reduces the danger of personal data misuse. The incorporation of the privacy legislation inside the framework ensures the legitimacy of any collection, processing or disclosure of personal data, while at the same time it relieves the service providers from the significant economical, operational and administrative overhead of keeping complex privacy policies up-to-date. Furthermore, by applying the respective policies, the presented framework can accurately handle the exemption cases of emergency situations, lawful interception and every other case where – for the common welfare – personal data revelation must be enforced.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

Current work focuses on assessing the presented framework in terms of performance and scalability and providing the means for addressing efficiently the respective issues raised by its deployment. Scalability refers to both establishing the means for load balancing between multiple Discreet Boxes that belong to a single organization and enabling the interoperation of Discreet Boxes for privacy-aware personal data exchange between organizations when necessary for advanced services provision. Additionally, the development of the necessary methodologies for the enhancement of the two ontologies is conducted, as well as the specification of the interfaces and guidelines for authoring new services and adapting existing ones in order to operate on top of middleware architecture and exploit its functionalities.

Acknowledgements

This work is partially supported by the European Union, in the framework of the FP6 – IST Project DISCREET [49]. The authors express their gratitude to the consortium for the fruitful discussions. Thanks are also due to the anonymous reviewers for their insightful comments.

References

[1] S.D. Warren, L.D. Brandeis, The right to privacy, Harvard Law Review IV (5) (1890) 193–220.
[2] The European Opinion Research Group, European Union citizens' views about privacy, Special Eurobarometer 196, December 2003.
[3] B.N. Schilit, N.I. Adams, R. Want, Context-Aware Computing Applications, in: Proceedings of Workshop on Mobile Computing Systems and Applications, Santa Cruz, CA, USA, December 1994.
[4] R. Cucchiara, Multimedia surveillance systems, in: Proceedings of the 3rd ACM International Workshop on Video Surveillance and Sensor Networks, Singapore, 2005, pp. 3–10.
[5] J. Han, M. Kamber, Data Mining: Concepts and Techniques, Morgan-Kaufman, NY, 2000.
[6] USA Federal Trade Commission, Eli Lilly Settles FTC Charges Concerning Security Breach, FTC File No. 012 3214, January 2002, http://www.ftc.gov/opa/2002/01/elililly.htm.
[7] USA Federal Trade Commission, FTC Sues Failed Website, Toysmart.com, for Deceptively Offering for Sale Personal Information of Website Visitors, FTC File No. 002 3274, July 2000, http://www.ftc.gov/opa/2000/07/toysmart.htm.
[8] European Parliament and Council, Directive 95/46/EC of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data, Official Journal of the European Communities, No. L 281, November 1995, pp. 31–50.
[9] E. Lundin, E. Jonsson, Anomaly-based intrusion detection: privacy concerns and other problems, Elsevier Computer Networks 34 (4) (2000) 623–640.
[10] United Nations, Universal Declaration of Human Rights, http://www.un.org/Overview/rights.html.
[11] US Public Law No. 93-579, December 31, 1974, 5 USC 552a.
[12] Organization for Economic Co-operation and Development, Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, September 1980.
[13] European Parliament and Council, Directive 2002/58/EC of the European Parliament and of the Council concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications), Official Journal of the European Communities, No. L 201, pp. 37–47, July 2002.
[14] European Parliament and Council, Directive 2006/24/EC of the European Parliament and of the Council of 15 March 2006 on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending Directive 2002/58/EC, Official Journal of the European Communities, No. L 105, April 2006, pp. 54–63.
[15] US Public Law No. 100-503, October 18, 1988, 5 USC 552a.
[16] S. Fischer-Hubner, IT – security and privacy, design and use of privacy-enhancing security mechanisms, Lecture Notes in Computer Science, vol. 1958, Springer, 2001.
[17] N.F. Noy, D.L. McGuinness, Ontology Development 101: A Guide to Creating Your First Ontology, Stanford Medical Informatics Technical Report SMI-2001-0880, March 2001.
[18] European Parliament and Council, Regulation 2195/2002/EC of the European parliament and of the council on the common procurement vocabulary (CPV), Official Journal of the European Communities, No. L 340, December 2002, pp. 1–562.
[19] Extensible Markup Language (XML) 1.0 (fourth ed.) Specification, W3C Recommendation, August 2004, http://www.w3.org/TR/2006/REC-xml-20060816.
[20] D. Crocker, P. Overel, Augmented BNF for syntax specifications: ABNF, IETF RFC 2234 (November) (1997).
[21] Hervé Äiache et al., Preliminary specification of the core architecture, IST DISCREET Deliverable D 2301 (November) (2006).
[22] A. Fuggetta, G.P. Picco, G. Vigna, Understanding code mobility, IEEE Transactions on Software Engineering 24 (5) (1998) 342–361.
[23] J. Claessens, B. Preneel, J. Vandewalle, (How) can mobile agents do secure electronic transactions on untrusted hosts? A survey of the security issues and the current solutions, ACM Transactions on Internet Technology 3 (1) (2003) 28–48.
[24] A. Westerinen et al., Terminology for policy-based management, IETF RFC 3198 (November) (2001).
[25] L.A. Sweeney, Computational Disclosure Control: Theory and Practice, Ph.D. Thesis, Massachusetts Institute of Technology, Boston, USA, 2001.
[26] E-ZPASS, New York Service Center, http://www.e-zpassny.com.
[27] R. Weinstein, RFID: a technical overview and its application to the enterprise, IEEE IT Professional 7 (3) (2005) 27–33.

[28] The Platform for Privacy Preferences 1.0 (P3P1.0) Specification, W3C Recommendation, April 2002, http://www.w3.org/TR/2002/REC-P3P-20020416/.
[29] A P3P Preference Exchange Language 1.0 (APPEL1.0), W3C Working Draft, April 2002, http://www.w3.org/TR/2002/WD-P3P-preferences-20020415.
[30] R. Agrawal, J. Kiernan, R. Srikant, Yirong Xu, XPref: a preference language for P3P, Elsevier Computer Networks 48 (5) (2005) 809–827.
[31] E. Bertino1, J. Byun1, N. Li, Privacy-preserving database systems, Foundations of Security Analysis and Design III, Lecture Notes in Computer Science, vol. 3655, Springer, 2005.
[32] G. Karjoth, M. Schunter, A Privacy Policy Model for Enterprises, in: Proceedings of the 15th IEEE Computer Foundations Workshop (CSFW '02), June 2002.
[33] G. Karjoth, M. Schunter, M. Waidner, Platform for enterprise privacy practices: privacy-enabled management of customer data, in: Proceedings of the 2nd Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 2482, Springer, 2002.
[34] M. Schunter, P. Ashley, The platform for enterprise privacy practices, in: Proceedings of the Information Security Solutions Europe (ISSE 2002), Paris, France, October 2002.
[35] G. Karjoth, M. Schunter, M. Waidner, Privacy-enabled Services for Enterprises, in: Proceedings of the International Workshop on Trust and Privacy in Digital Business (TrustBus 2002), Aix en Provence, France, September 2002.
[36] P. Ashley, S. Hada, G. Karjoth, C. Powers, M. Schunter, The Enterprise Privacy Authorization Language (EPAL), EPAL 1.2 Specification, IBM Research Report, 2003, http://www.zurich.ibm.com/security/enterprise-privacy/epal/.
[37] M. Backes, B. Pfitzmann, M. Schunter, A toolkit for managing enterprise privacy policies, in: Proceedings of the 8th European Symposium on Research in Computer Security, Gjvik, Norway, October 13–15, Lecture Notes in Computer Science, vol. 2808, Springer, 2003.
[38] Organization for the Advancement of Structured Information Standards, OASIS eXtensible Access Control Markup Language (XACML) TC, 2004, http://www.oasis-open.org/committees/xacml/.
[39] M. Casassa Mont, S. Pearson, P. Bramhall, Towards accountable management of privacy and identity information, in: Proceedings of the 8th European Symposium on Research in Computer Security (ESORICS 2003), Gjøvik, Norway, October 2003.
[40] M. Casassa Mont, R. Thyne, P. Bramhall, Privacy Enforcement with HP Select Access for Regulatory Compliance, Hewlett-Packard Labs Technical Report, HPL-2005-10, 2005.
[41] M, Casassa Mont, Dealing with privacy obligations: important aspects and technical approaches, in: Proceedings of the International Workshop on Trust and Privacy in Digital Business (TrustBus 2004), Zaragoza, Spain, August 2004.
[42] M. Casassa Mont, R. Thyne, A systemic approach to automate privacy policy enforcement in enterprises, in: Proceedings of the 6th Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 4258, Springer, 2006.
[43] D.F. Ferraiolo, R. Sandhu, S. Gavrila, R.D. Kuhn, R. Chandramouli, Proposed NIST standard for role-based access control, ACM Transaction Information and System Security 4 (3) (2001).
[44] IST PRIME (Privacy and Identity Management for Europe) project, home page: https://www.prime-project.eu/.
[45] IST FIDIS (Future of Identity in the Information Society) project, home page: http://www.fidis.net/.
[46] The Liberty Alliance Project, home page: http://www.projectliberty.org/.
[47] B. Pfitzmann, Privacy in enterprise identity federation, in: Proceedings of the 3rd Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 2760, Springer, 2003.
[48] M. Alsaleh, C. Adams, Enhancing consumer privacy in the liberty alliance identity federation and web services frameworks, in: Proceedings of the 6th Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, vol. 4258, Springer, 2006.
[49] IST Project DISCREET (Discreet Service Provision in Smart Environments), home page: http://www.ist-discreet.org.

Nikolaos L. Dellas was born in Vonitsa, Greece, on December 10, 1980. He received his Dipl.-Ing. degree from the National Technical University of Athens (NTUA) in 2003 in Electrical and Computer Engineering. Since then, he is a research associate and a Ph.D. candidate in the Intelligent Communications and Broadband Networks Laboratory of NTUA. His primary research interests are security and privacy solutions, mobile computing, middleware technologies and performance evaluation of distributed technologies. He has participated in several European and National research projects. He is a member of the Technical Chamber of Greece.

Dimitra I. Kaklamani was born in Athens, Greece, in 1965. She received the Diploma and Ph.D. degrees from the School of Electrical and Computer Engineering National Technical University of Athens (NTUA), in 1987 and 1992, respectively. From 1993–1994, she was a researcher at the Institute of Communications and Computer Systems, NTUA. From January 1994 to September 1996, she was a Consultant at the Hellenic Telecommunications Organization S.A. In April 1995, May 2000 and October 2004 she was elected Lecturer, Assistant Professor and Associate Professor, respectively, at the School of Electrical and Computer Engineering, NTUA. She is the coeditor of Applied Computational Electromagnetics, State of the Art and Future Trends (NATO ASI Series F: Computer and System Sciences, Vol. 171, New York: Springer-Verlag, 2000) and has published over 150 journal and conference papers. Her current research interests focus on the use of object oriented methodologies and middleware technologies for the development of distributed systems, as well as development of visualization and real-time simulation techniques for solving complex large-scale modeling problems of microwave engineering and information transmission systems. She is a member of the Technical Chamber of Greece.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010

Sofia H. Kapellaki was born in Athens, Greece on April 13, 1977. She received her Dipl.-Ing. degree from the Mechanical Engineering and Aeronautics Department of the University of Patras, Patras, Greece in 2000 and her Ph.D. from the School of Electrical and Computer Engineering of the National Technical University of Athens (NTUA) in 2007. Her primary research interests are telecommunications software design and implementation, service engineering and open network architectures. She has also experience on service discovery, distributed objects and middleware technologies. She has participated in several National and European Union research projects. She is a member of the Technical Chamber of Greece.

Eleftherios A. Koutsoloukas was born in Athens, Greece, on August 10, 1979. He received his Dipl-Ing. degree from the Electrical and Computer Engineering Department of the National Technical University of Athens (NTUA) in 2002. Since then, he is a research associate and a Ph.D. candidate in the Intelligent Communications & Broadband Networks Laboratory of NTUA. His primary research interests are mobile computing, software development in the area of mobile Internet, middleware applications and distributed technologies. He is a member of the Technical Chamber of Greece.

Georgios V. Lioudakis received his Dipl.-Ing. degree in Electrical and Computer Engineering from the National Technical University of Athens (NTUA) in 2002. Since then, he is a research associate of the Intelligent Communications and Broadband Networks Laboratory of NTUA, while he's working towards the Ph.D. degree in the area of privacy protection. His research interests lie mainly in privacy protection, mobile networks and middleware and distributed technologies. He has participated in several European and National research projects and has published a number of papers related to the above fields. He is a member of the IEEE, the ACM and the Technical Chamber of Greece.

George N. Prezerakos was born in Athens, Greece, in December 1970. He earned the Dipl.-Ing. from the National Technical University of Athens (NTUA), in 1993 in electrical and computer engineering and the Ph.D. in 1998 also from NTUA in the area of Broadband Networks. From 1993 to 1998 he was a research associate of the NTUA Telecommunications Laboratory performing research in the area of Intelligent Network, Multimedia Service Engineering for Broadband Networks and Fuzzy Logic Systems. He is currently an Associate Professor in the Electronic Computing Systems Department of the Technological Education Institute (TEI) of Piraeus. His current research interests are in the fields of service and platform engineering, context-aware services and model driven development. He has several publications in the above areas. He has participated in several European Union and national projects. He is member of ACM and the Technical Chamber of Greece.

Nikolaos D. Tselikas was born in Athens, Greece, on May 14, 1976. He received both his Dipl.-Ing. degree and his Ph.D. from the School of Electrical and Computer Engineering of the National Technical University of Athens (NTUA) in 1999 and 2004, respectively. He is currently working in NTUA as a research associate. He is interested in broadband Intelligent Networks, network convergence, service control, distributed and middleware architectures and telecommunications service engineering. He has participated in several European Union and National research projects. He is a member of the Technical Chamber of Greece.

Iakovos S. Venieris was born in Naxos, Greece, on March 3, 1965. He received the Dipl.-Ing. degree from the University of Patras, Patras, Greece in 1988, and the Ph.D. degree from the National Technical University of Athens (NTUA), Athens, Greece, in 1990, all in electrical and computer engineering. In 1994 he became an Assistant Professor in the School of Electrical and Computer Engineering of NTUA, where he is now a Full Professor. His research interests are in the fields of broadband communications, Internet, mobile networks, Intelligent Networks, internetworking, signalling, service creation and control, distributed processing, privacy protection, agents technology, and performance evaluation. He has over 150 publications in the above areas. He leads NTUA participation in several European Union and National Projects. He is a reviewer for several IEEE, ACM, Elsevier, and John Wiley journals, associate editor for the IEEE Communications Letters, member of the editorial staff of Computer Communications (Elsevier), and has been guest editor in the IEEE Communications Magazine. He is the co-editor of two international books on Intelligent Broadband Networks (Wiley, 1998) and Object oriented Software Technologies in Telecommunications (Wiley, 2000). He is a member of IEEE and the Technical Chamber of Greece.

Please cite this article in press as: G.V. Lioudakis et al., A middleware architecture for privacy protection, Comput. Netw. (2007), doi:10.1016/j.comnet.2007.06.010