The Architecture of a Privacy-aware Access Control Decision Component

Claudio A. Ardagna, Marco Cremonini, Ernesto Damiani,
Sabrina De Capitani di Vimercati and Pierangela Samarati

Dipartimento di Tecnologie dell'Informazione,
Università degli Studi di Milano
26013 Crema, Italy
{ardagna,cremonini,damiani,decapita,samarati}@dti.unimi.it

Abstract. Today many interactions are carried out online through Web sites and e-services and often private and/or sensitive information is required by service providers. A growing concern related to this widespread diffusion of on-line applications that collect personal information is that users' privacy is often poorly managed and sometimes abused. For instance, it is well known how personal information is often disclosed to third parties without the consent of legitimate data owners or that there are professional services specialized on gathering and correlating data from heterogeneous repositories, which permit to build user profiles and possibly to disclose sensitive information not voluntarily released by their owners. For these reasons, it has gained great importance to design systems able to fully preserve information privacy by managing in a trustworthy and responsible way all *identity and profile information*.
In this paper, we investigate some problems concerning identity management for e-services and present the architecture of the Access Control Decision Function, a software component in charge of managing access request in a privacy-aware fashion. The content of this paper is a result of our ongoing activity in the framework of the PRIME project (*Privacy and Identity Management for Europe*) [18], funded by the European Commission, whose objective is the development of privacy-aware solutions for enforcing security.

1 Introduction

From the growing offering of e-services provided by a number of organizations, users have not only gained benefits in terms of variety and richness of accessible services. The drawback of such an increase in service provision is that a corresponding growing amount of personal information is communicated by users of e-services to the corresponding providers. Personal identifiable information (PII) are required by e-service providers for many legitimate reasons (e.g., to offer personalized services). Also, requiring personal information permits to mitigate abuses of e-services and to

avoid, for example, the access by means of automatic software instead of physical users. Finally, personal information of e-service users is needed for marketing purposes, such as promoting new services or producing access statistics for advertisers.

However, despite all these reasons for collecting personal information are certainly legitimate, many concerns exist about the privacy of e-service users. Such concerns are motivated by observing that the number and type of personal information collected by service providers permit to easily profile user's habits and preferences in a very detailed and precise way. In addition, it is well known how personal information is often disclosed to third parties without the consent of legitimate data owners or that there are professional services specialized on gathering and correlating data from heterogeneous repositories, which permit to build user profiles and possibly to disclose sensitive information not voluntarily released by their owners.

As a consequence, users concerned about their private information are increasingly refusing to benefit from such a widespread offering of e-services because they prefer not to have their personal data under the control of anyone at anytime.

A key aspect to address these concerns is the notion of *privacy-aware access control*, which encompasses and combine the notions of *privacy* and of *access control* in an homogeneous framework. Traditional access control systems are based on regulations (policies) that establish who can, or cannot, execute certain actions on some resources and the way they compute access decisions is based on the requester's credentials carrying her identity and other personal information (e.g., affiliation, membership, and so on) [10].

Other requirements that traditional access control systems usually do not take into account are related to *data usage*, which is the possibility to specify how data accessed by an authorized party must be handled. This represents a novel feature for access control that is no simply concerned with authorizing the access to data and resources but also with defining and enforcing the way data and resources are subsequently managed. Also, in modern systems, the definition of an access control model is complicated by the need to formally represent complex policies, where access decisions depend on the application of different rules coming from laws practices, organizational regulations, and so on.

Privacy awareness and features to manage requesters credentials accordingly are not taken into account by access control systems in use today. Requiring privacy awareness means that credentials and personal

information of users that request e-services cannot be freely available and manageable by service providers. Privacy poses constraints on which data can be required for a certain service and on the way personal information once collected by a service provider can be handled, released to third parties, or recorded.

Despite recent advancements in access control models have permitted to use generic attributes/properties of both requesters and resources, access control systems are not yet designed for enforcing privacy policies.

Therefore, by considering privacy issues, there is the need to improve authorization policies and models and to develop new solutions for access control, authorization specification, and enforcement. The development of such solutions will require to investigate open research problems as well as to implement an access control architecture addressing privacy concerns from its foundations.

In this paper, we describe an approach aimed at providing users with a privacy-aware access control system that enforces privacy requirements. In particular, we present the architecture of the *Access Control Decision Function* (ACDF), an autonomous software component for controlling access to data in the framework of e-services. The ACDF component is based on a flexible model and XML-based language [2]. Our work has been carried out in the context of the Privacy and Identity Management for Europe (PRIME) project, an European project whose goal is the development of privacy-aware solutions for enforcing security.

The remainder of this paper is organized as follows. Section 2 summarizes the main contributions in the field of privacy-aware access control and describes the way our approach differs from the previous ones. Section 3 describes the new requirements for a privacy-aware access control and gives an overview of the PRIME project. Section 4 summarizes our proposal for a privacy-aware access control policy. Section 5 presents the architecture of the Access Control Decision Function, explaining its interactions with external components and the overall work flow. Finally, Section 6 draws our conclusions and sketches future work.

2 Related Work

A number of projects and research papers about privacy have been presented in the last few years, although not many of them have addressed the issue of privacy-aware access control. More in detail, two lines of research are closely related to the topic of this paper: *i)* the definition and development of access control and privacy languages, and *ii)* the defini-

tion of infrastructures to protect and preserve privacy of either services or clients.

For what concerns the first research topic, some languages have been defined starting from languages for access control as *XACML* (eXtensible Access Control Markup Language) [22] to data handling languages (i.e., languages regulating how personal information could be managed once collected) as for instance *P3P* (Platform for Privacy Preferences Project) [5, 8] and *EPAL* (Enterprise Privacy Authorization Language) [4, 5].

XACML [22] is an XML-based language used to define access control policies. The main differences between XACML and the language developed for our ACDF component are that XACML does not consider data handling constraints, it does not explicitly support neither privacy features nor variables in the definition of policies (a feature that permits to greatly enhance policy expressiveness), and it is not integrated with the ontological approach that our ACDF solution exploits in the more general context of the PRIME Project. In addition to the language, XACML defines both an architecture for the evaluation of policies and a communication protocol for messages interchange. The most important difference between the XACML's system design and architecture and our proposal is that XACML assumes to have all the information about a requester available at the time of policy evaluation and access control decision. In our ACDF component, instead, a negotiation phase between a requester and a provider is carried out in order to establish the number and type of credentials that, on the one hand, are sufficient for the service provision and, on the other hand, minimize the disclosure of personal information.

P3P [5, 8] is a project widely acknowledged that addresses the need of a user to assess that the privacy practices adopted by a server provider comply with her privacy requirements. Supporting data handling policies in Web-based transactions is the goal of P3P, which permits the definition of server privacy practices in a standard format, allowing users to automatically understand and match these practices against their privacy preferences. Thus, users need not read the privacy policies at every site they interact with but they are always aware of the server practices in data handling. Some drawbacks of P3P are the lacking of a formal and unambiguous language to define user privacy preferences, of a technical mechanism to verify that Web sites respect users policies and of a process to negotiate the privacy practices between the interacting parties. In addition, P3P scope is restricted to Web sites only.

EPAL [4, 5] is an XML-based markup language that formalizes enterprise-internal privacy policies. It approaches the problem on the server side and addresses the need of a company to specify access control policies, with reference to attributes/properties of the requestor, to protect private information of its users. EPAL is designed to enable organizations to translate their privacy policies into IT control statements and to enforce policies that may be declared and communicated in P3P. XACML, however, includes most (if not all) of the expressive power of EPAL.

Considering projects that aim at developing an architecture to preserve security and privacy, several have been proposed. International Security, Trust, and Privacy Alliance (ISTPA) [13] is an open, policy-configurable model consisting of several privacy services and capabilities, intended to be used as a template for designing solutions and covering security, trust, and privacy requirements. The goal of the framework is to set the basis for developing products and services that support current and evolving privacy regulations and business policies.

Reasoning on the Web with Rules and Semantics (REWERSE) [6, 19] is an european network of excellence on the semantic web whose objective is to enrich the Web with so-called intelligent capabilities for data and service retrieval, composition, and processing. REWERSE's research activities will be devoted to several objectives such as *policy specification, composition, and conformance* aiming at user-friendly high-level specifications for complex Web systems.

Enterprise Privacy Architecture (EPA) [17] is an IBM project that wants to improve enterprises e-business trust. EPA represents a new approach to privacy that tries to help organizations to understand how privacy impacts business processes. EPA defines privacy parties, rules, and data for new and existing business processes and provides privacy management controls based on consumer preferences, privacy best practices, and business requirements.

Finally, TRUSTe [21] is an organization dedicated to preserving customer privacy and assisting e-commerce with customer privacy concerns. It certifies and monitors Web site privacy practices.

3 Requirements for a privacy-aware access control

In general, an environment well-suited for users needing a private and secure way for using e-services should support at least the following basic requirements.

- *Privacy.* A digital identity solution should be respectful of the users rights to privacy and should not disclose personal information without explicit consent.
- *Minimal disclosure.* Service providers must require the least set of credentials needed for service provision, and users should be able to provide credentials selectively, according to the type of on-line services they wish to access.
- *Anonymity support.* As a special but notable case of minimal disclosure, many services do not need to know the real identity of a user. Pseudonyms, multiple digital identities, and even anonymous accesses must be adopted when possible.
- *Legislation support.* Privacy-related legislation is becoming a powerful driver toward the adoption of digital identities. The exchange of identity data should not then violate government legislation such as the Health Insurance Portability and Accountability Act (HIPPA) or Gramm-Leach-Bliley Act (GLB).

With respect to these privacy-based requirements, the usual way of designing access control systems is not satisfactory. In particular, selective disclosure of credentials is normally not implemented, because users' attributes, for example inserted into X.509 identity certificates [14] or collected as attribute certificates [11], are defined according to functional needs, making it easier to collect all credentials in a row instead of iteratively asking for the ones strictly necessary for a given service only. With XACML the same requirement holds and credentials are collected entirely before policy evaluation. Pseudonymity, multiple identities, and anonymity are also usually not supported.

These new requirements regarding an improved management of digital identities are among the motivations of the PRIME project [18], a large-scale research effort aimed at developing an identity management system able to protect users personal information and to provide a framework that can be smoothly integrated with current architectures and on-line services.

More specifically, providing the users with the control of their personal data and permitting anonymous interactions are some of the main goals of the PRIME project. Next, users should also be able to use different pseudonyms during interactions with other parties, a feature that reduces the risk of profiling by making different transactions performed by the same user unlinkable one with the others. Another goal of the PRIME project is to define privacy rules governing the system usage. The rules should establish how to use the system and, in particular, allow

the definition of policies to define trust relationships, privacy preferences, and authorization rules.

Following the definition of an enhanced authorization model based on privacy awareness, policies must be effectively enforced at the receiving end. The enforcement of privacy policies is a more complicate task than the enforcing of traditional access control policies because they have several additional features such as obligations, policy composition and negotiation. The privacy-enhancing technical components developed within the PRIME project will be integrated to produce a privacy-enhancing digital identity management system [1, 3, 15].

4 A privacy-aware access control model and language

To define a privacy-enhanced access control system based on the concept of digital identity, we first need to identify the main characteristics that the corresponding access control model should possess.

- *Policy formats.* Parties need to specify protection requirements on the data they make available using a format both human and machine readable, easy to inspect and interchange.
- *Access control rules.* Access control rules should be able to make use of partial identities associated with users. Also, it is important to be able to specify access control rules about subjects accessing the information and about resources to be accessed in terms of rich ontology-based metadata (e.g., Semantic Web-style ones) increasingly available in advanced e-service applications [9].
- *User-driven constraints.* In addition to traditional server-side access control rules, users should be able to specify constraints and restrictions about the usage that will be done of their information once released to a third party.
- *Interactive enforcement*. A novelty of our framework is that we do not assume anymore that all credentials are collected before an access request is evaluated. Instead, the access control component may not have all the information it needs to decide whether or not an access should be granted. On the other side, the requester may not know in advance which information will be asked to get the access to the service. As a consequence, a new way of enforcing the access control process has been defined based on a negotiation protocol aimed at establishing the least set of information that the requester has to disclose in order to access the desired service.

To take all these issues into account, a new privacy-aware access control model together with an access control protocol for the communication of policies and of identity information among parties have been defined and the following different types of privacy policies have been introduced:

- traditional *access control policies* governing access/release of data/services managed by the party [20];
- *release policies* governing the release of properties, credentials, and personal identifiable information of the party [7];
- *data handling policies* defining how personal information released by a third party have to be managed [8];
- *sanitized policies* filtering the response to be returned to the counterpart to avoid release of sensitive information related to the policy itself.

In the following, we focus on access control and release policies.

4.1 Privacy-aware access control rules

Although it is not in the scope of this paper to discuss the details of the access control language, a brief introduction of its basic elements is necessary to describe the different sub-systems that must be coordinated together with the ACDF. In short, the main elements of PRIME's authorization rules are as follows.

- *Subject expression*: a boolean formula of terms that allows the reference to a set of subjects depending on whether they satisfy or not certain conditions, where conditions can evaluate the user's profile, location predicates, or the user's membership in groups, roles, and so on.
- *Object expression*: a boolean formula of terms that allows the reference to a set of objects depending on whether they satisfy or not certain conditions, where conditions evaluate membership of the object in categories, values of properties on metadata, and so on.
- *Actions*: the action (or class of actions) to which the rule refers.
- *Purposes*: a statement, certified or not, representing how the data is going to be used by the recipient.
- *Conditions*: a boolean formula of terms that express additional conditions, for example, dictated by legislation, location-based conditions, and trust conditions.

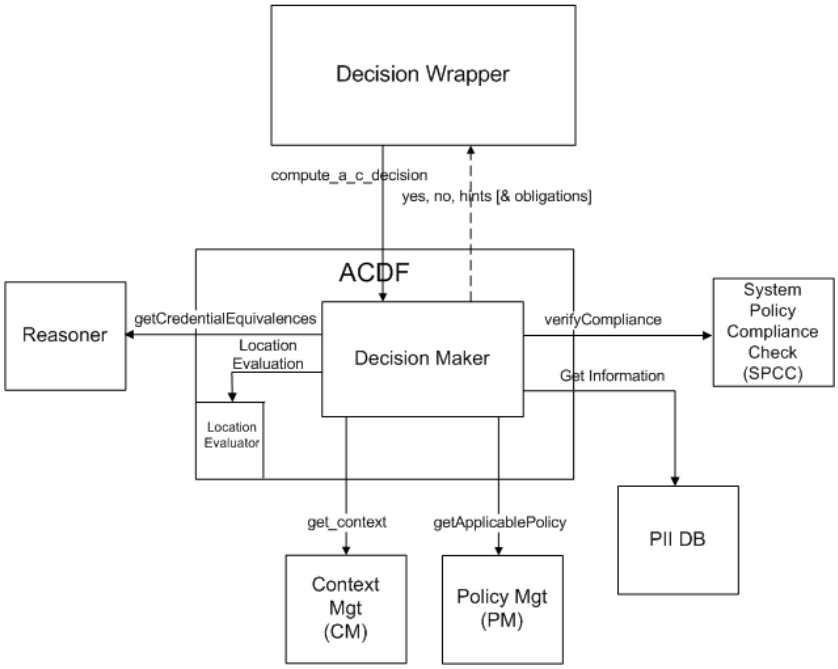

Fig. 1. The Access Control Decision Function and its interactions with other components.

– *Obligations*: conditions defined by the users and attached to corresponding data when they are disclosed to third parties. Receiving parties must comply with obligations coming along with data and the framework is able to enforce it.

Each access request results in an *access decision* that can take three different forms:

– *Yes*: the access request is granted;
– *No*: the access request is denied;
– *Undefined*: the access request provides insufficient information to determine whether the request can be granted or denied. The negotiation phase between the requester and the service provider is entered.

5 ACDF architecture

The PRIME's Access Control component is composed by two parts: the Access Control Decision Function (ACDF) responsible for taking an access decision for all access requests directed to PRIME resources, like data and services, and the Access Control Enforcement Function (ACEF) responsible for the enforcing of access control decisions by intercepting accesses to resources and granting them only if they are part of an operation for which a positive decision has been taken. From an architectural point of view, the ACDF is a unique module composed by different sub-modules associated with specific tasks of the decisional process or in charge of interacting with external components. More precisely, the submodules are the following.

- *Decision Maker*: produces the final response possibly combining different access decisions coming from different sub-components;
- *Policy Evaluator*: manages the evaluation of the applicable policies against an access request;
- *Policy Handler*: is in charge of managing all communications with the Policy Manager (an external component) to retrieve all policies applicable to an access request;
- *Reasoner Administrator*: manages communication with the Reasoner component to require reasoning operations about policies to calculate extended policies;
- *Context Administrator*: manages the access and the communication with the Context Manager component, which is the requestors information repository during a transaction;
- *PII Database Mediator*: manages the communication with the information (PII) repository that represents the storage system for personal information;
- *SPCC Handler*: manages all interactions with the System Policy Compliance Check (SPCC) component, the one in charge of evaluating special conditions based on assurance and trust predicates;
- *LBS Evaluator*: is the sub-module that evaluates special conditions based on location-based predicates;
- *Obligation Handler*: selects and attaches to the access decision all corresponding obligations.

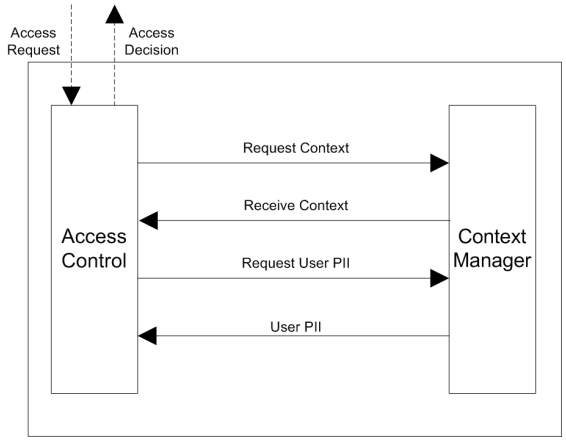

Fig. 2. Interactions with the Context Manager

5.1 ACDF Interactions

As illustrated in Figure 1, the ACDF component interacts with many other components of PRIME's Identity Management System (IDMS). Below we present a brief description of these components.

Context Manager (CM). The Context Manager component manages user's *session data* (see Figure 2). Session thereby denotes a single communication action, usually one connection established by an access requester. The context management acts as a database for the ACDF that can query it for retrieving credentials (*User PII*).

The data structure of a single context contains information on the following two aspects:

- data disclosed to and by the communication party such as pseudonym and personal information, either certified or not;
- certified proofs about negotiation, disclosure, and exchange of personal information.

Policy Manager (PM). The Policy Manager component handles the life cycle management of policies by providing functionalities for policy administration (see Figure 3). Related to the access decision, the ACDF interacts with the Policy Manager to collect all policies that can be applied to the access request being evaluated. The Policy Manager has a searching functionality that filters out policies based on access request attributes.

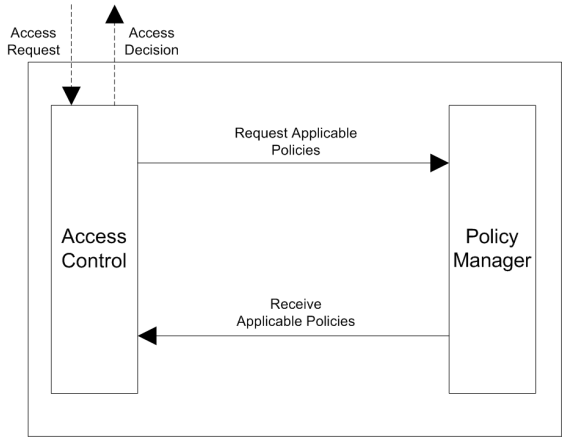

Fig. 3. Interactions with the Policy Manager

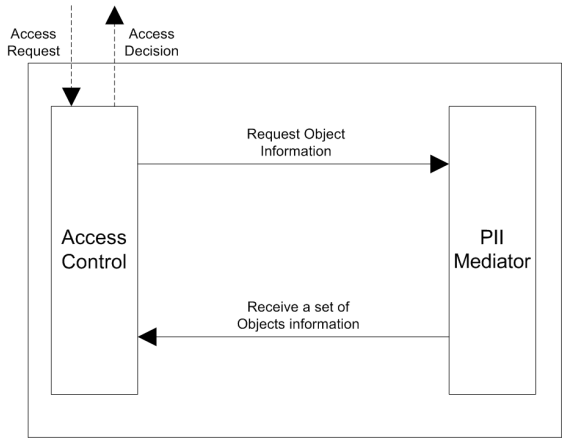

Fig. 4. Interactions with the PII Database Mediator

PII Database Mediator (PII DB). The PII Database Mediator component manages all accesses to the database containing personal information (*PII*) (see Figure 4). The access to PII information stored into the PII Database is handled by the Mediator component so that no special privilege is granted to internal modules of PRIME. The ACDF interacts with the PII Database Mediator by invoking a specific method and passing all parameters needed for querying PII data.

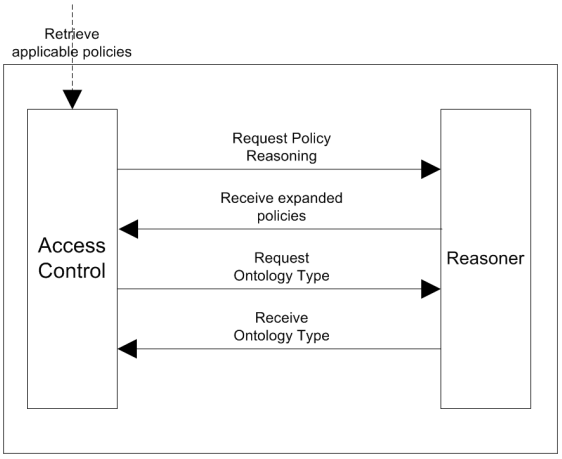

Fig. 5. Interactions with the Reasoner

Reasoner. The Reasoner is the component that maintains and makes use of the ontologies defined in the project (see Figure 5). It provides deductions based on machine readable data and rules. In addition to data and prolog style rules and generic methods for producing all inferences from an ontology, the module also provides methods specific to some PRIME components. For the ACDF component, in particular, this includes *credentials equivalences*, which is a feature to verify equivalences between credential expressed according to different ontologies. The reasoner is based on the Jena API and as such requires data and ontologies to be expressed using Jena RDF models [16].

System Policy Compliance Check (SPCC). The SPCC module handles trust, assurance and accountability compliance conditions which requires the analysis of the assurance information (see Figure 6). Although in certain cases trust and assurance constraints, specified by some policies, can be computed statically and independently of access control, in other cases (notably when dynamic constraints are involved) trust conditions need to be evaluated together with other conditions by the ACDF. In these cases, the ACDF recognizes the assurance constraint during the evaluation process and invokes the SPCC component to evaluate it.

5.2 Decision Maker

Having introduced all the components involved in the access control process, the core module of ACDF, the *Decision Maker*, can be fully de-

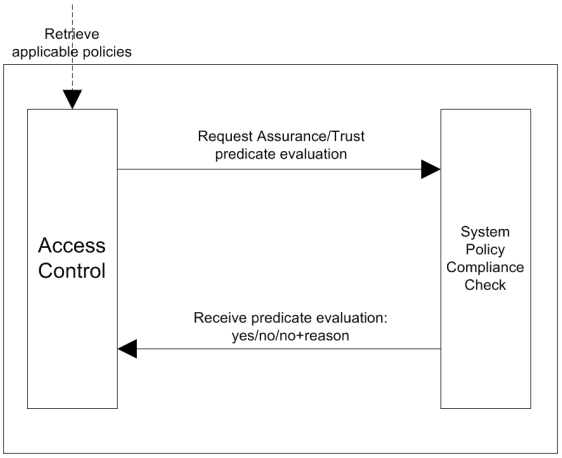

Fig. 6. Interactions with the SPCC

scribed. The Decision Maker is the module responsible for all access control decisions and returns a *Yes*, *No*, or *Undefined* response. It handles applicable access control policies and proceeds evaluating all different components of the rules, like subject expressions, object expressions, and so on. Such an evaluation requires the Decision Maker to interact and coordinate with both sub-modules internal to the ACDF and external components. The ACDF execution flow prescribes that, first, the ACDF receives an access request and selects the context associated with the current session through a *Context Manager* API. After that, information related to the requested object is collected from the *PII Database Mediator* and all the applicable policies are retrieved from the *Policy Management* module by means of access request attributes. When applicable policies are acquired, the evaluation process can start and proceeds as follows:

1. predicates based on trust/assurance properties are communicated to the *System Policy Compliance Check (SPCC)* that is in charge of evaluating them;
2. predicates about the subject are evaluated based on context information;
3. similarly, predicates about the object of the request are evaluated by interacting with the PII Database Mediator;
4. location-based predicates represent a special case and their evaluation is delegated to a specialized sub-module of ACDF, called *Location Evaluator*;

5. with all partial evaluations generated by sub-modules and external systems, the Decision Manager produces a final access control decision by composing all partial decisions:
   (a) if the decision is *positive* (response *Yes*) obligations and constraints need to be returned. Obligation defines how released data must be handled after disclosure, constraints provide directives to the PII Database Mediator when the access is enforced;
   (b) if the decision is *negative* (response *No*) a reason for that can be returned attached to the answer;
   (c) if a decision cannot be reached (response *Undefined*) obligations and additional requests are returned to the subject, possibly sanitized for preventing disclosure of access control policies details.

Finally, the ACDF produces a message for the *Decision Wrapper*, which acts as a mediator between the ACDF and the ACEF module, to communicate to the ACEF component how to handle the corresponding access request.

As an example, consider a rent-a-car scenario and suppose that a policy states that "an anonymous user with a valid Italian driver license can rent a sport car with a special price of 80 euro per day, if she is in Italy, she is more than 21 years old, and if the rent-a-car service provider has a working trusted platform management".

Figure 7 illustrates a representation of this policy using our privacy-aware access control language.

At server-side, suppose now that an access request stating that "Mary want to rent a sport car" arrives together with her credentials. The ACDF can query the context from the Context Manager to verify the age, the location, and the availability of an ecoin card of Mary. Assume that, among the required credentials, the driver license is missing.

The Decision Maker decomposes the policy and sends the location-based predicate (`lbs` element) to the LBS Evaluator, the assurance predicate (`trust` element) to the SPCC Handler, and evaluates the remaining conditions. After the evaluation, the LBS evaluator returns a positive response (Mary is in Italy), the SPCC handler returns a positive response (the server has a working TPM), and the Policy Evaluator calculates an undefined response due to the fact that Mary has not previously released her driver license. The Decision Maker collects all these responses and returns a final undefined decision together with a request for the driver license. At this point, Mary based on her privacy preferences can decide whether to disclose her driver license or to terminate the transaction.

```
<policy>
 <subject>any</subject>
 <action>http://.../action#rent_a_car</action>
 <object>http://.../object#sport_car</object>
 <purpose>http://.../purpose#any_purpose</purpose>
 <subjectExpr>
  <condition>
   <Lval>Idemix-EU-DriversLicence.Issuer.Country</Lval>
   <op>=</op>
   <rVal>IT</rVal>
  </condition>
  <condition>
   <Lval>
    Idemix-EU-DriversLicence.Permit.CarPermit.Allowed
   </Lval>
   <op>=</op>
   <rVal>true</rVal>
  </condition>
  <condition>
   <Lval>User.Age</Lval>
   <op>></op>
   <rVal>21</rVal>
  </condition>
 </subjectExpr>
 <objectExpr/>
 <trust>http://.../assurance#HasWorkingTMP</trust>
 <lbs>in_area("Italy")</lbs>
 <genCond>
  <condition>
   <Lval>Idemix-Ecoin.Value</Lval>
   <op>=</op>
   <rVal>80</rVal>
  </condition>
 </genCond>
 <ns>http://.../prime-PII-lite</ns>
 <obligation ref="OBL1">
</policy>
```

Fig. 7. A simple example of policy.

6 Conclusions and Future Work

To protect the privacy of parties in today's global infrastructure we need to combine solutions from technology, legislation, and organizational practices. This paper showed a first proposal towards the solution of this problem developed in the context of our ongoing activity in the framework of the PRIME project. In particular, with respect to previous privacy-

aware access control frameworks, this solution fully takes into account the possibility for the user to negotiate the credentials to be released and actually permits to enforce the principle of minimal disclosure. The solution, moreover, is not strictly targeted to Web-based transactions and to data handling policies, as for P3P. Future work include the development of negotiation policies to be applied to the parties; the extension of the notion of subject ontology to capture more complex assertions on subjects, as well as the notion of object and credential ontology; the support of variables into the language to achieve a higher degree of expressiveness.

7 Acknowledgments

This work was supported in part by the European Union within the PRIME Project in the FP6/IST Programme under contract IST-2002-507591 and by the Italian MIUR within the KIWI and MAPS projects.

References

1. C.A. Ardagna, E. Damiani, S. De Capitani di Vimercati, P. Samarati. A Web Service Architecture for Enforcing Access Control Policies. *In Proc. of the First International Workshop on Views On Designing Complex Architectures (VODCA 2004)*, Bertinoro, Italy, September 11-12, 2004.
2. C.A. Ardagna, E. Damiani, S. De Capitani di Vimercati and P. Samarati. Towards Privacy-Enhanced Authorization Policies and Languages. In *Proc. of the 19th Annual IFIP WG 11.3 Working Conference on Data and Applications Security (IFIP)*, Nathan Hale Inn, University of Connecticut, Storrs, USA, August 7-10, 2005.
3. C.A. Ardagna and S. De Capitani di Vimercati. A comparison of modeling strategies in defining XML-based access control languages. *Computer Systems Science & Engineering Journal*, 2004.
4. P. Ashley, S. Hada, C. Powers and M. Schunter. Enterprise Privacy Authorization Language(EPAL). *IBM Research*, 2003.
5. P. Ashley, S. Hada, G. Karjoth and M. Schunter. E-P3P privacy policies and privacy authorization. *In Proc. of the ACM workshop on Privacy in the Electronic Society (WPES 2002)*, Washington, DC, USA, November 21, 2002.
6. P. A. Bonatti and D. Olmedilla. Driving and monitoring provisional trust negotiation with metapolicies. *In Proc. of the IEEE 6th International Workshop on Policies for Distributed Systems and Networks (POLICY 2005)*, Stockholm, Sweden, 6-8 June 2005.
7. P. Bonatti and P. Samarati. A unified framework for regulating access and information release on the web. *Journal of Computer Security*, 10(3):241–272, 2002.
8. L. Cranor and M. Langheinrich and M. Marchiori and M. Presler-Marshall and J. Reagle. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. `http://www.w3.org/TR/P3P/`.

9. E. Damiani, A. Corallo, G. Elia. A Knowledge Management System Enabling Regional Innovation. *In Proc. of the VI international conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2002)*, Crema, Italy, September 16-18 2002.
10. S. De Capitani di Vimercati, S. Paraboschi, and P. Samarati. Access control: Principles and solutions. *Software – Practice and Experience*, 33(5):397–421, April 2003.
11. S. Farrell and R. Housley. An Internet Attribute Certificate for Authorization. Request For Comments 3281, *Internet Engineering Task Force*, 2002.
12. C. A. Gunter, M. J. May and S. G. Stubblebine. A Formal Privacy System and its Application to Location Based Services. *In Proc. of the 4th Workshop on Privacy Enhancing Technologies (PET 2004)*, Toronto, Canada, May 26-28, 2004.
13. International Security, Trust, and Privacy Alliance (ISTPA), `http://www.istpa.org/`
14. ITU Telecommunication Standardization Sector (ITU-T). Information Technology - Open Systems Interconnection - The Directory: Authentication Framework. Recommendation X.509 (03/00), *International Telecommunication Union*, 2000.
15. S. Jajodia, P. Samarati, M. Sapino, and V. Subrahmanian. Flexible support for multiple access control policies. *ACM Transactions on Database Systems*, 26(2):18–28, June 2001.
16. Jena. `http://jena.sourceforge.net`.
17. G. Karjoth, M. Schunter and M. Waidner. Privacy-enabled Services for Enterprises *In Proc. of the 13th International Conference on Database and Expert Systems Applications (DEXA'02)*, Aix-en-Provence, France, September 2-6, 2002.
18. PRIME (Privacy and Identity Management for Europe). `http://www.prime-project.eu.org`.
19. Reasoning on the Web (REWERSE), `http://www.pms.ifi.lmu.de/rewerse-wga1/index.html`
20. P. Samarati and S. De Capitani di Vimercati. Access control: Policies, models, and mechanisms. In R. Focardi and R. Gorrieri, editors, *Foundations of Security Analysis and Design*, LNCS 2171. Springer-Verlag, 2001.
21. Truste, `http://www.truste.org/about/index.php`
22. XACML - (eXtensible Access Control Markup Language), `http://www.oasis-open.org/committees/tc\_home.php?wg\_abbrev=xacml\#XACML20`