TOWARDS IDENTITY MANAGEMENT FOR E-SERVICES

Claudio A. Ardagna,[1] Marco Cremonini,[2] Ernesto Damiani,[3]
Sabrina De Capitani di Vimercati,[4] Pierangela Samarati[5]

Abstract

Nowadays, business and recreational activities are increasingly based on the use of remote resources and e-services, and on the interaction between different, remotely located parties. In such a context, it is of paramount importance that electronic execution of private and/or sensitive transactions fully preserves information privacy, managing in a trustworthy and responsible way all *identity and profile information* to be released to remote parties. In this paper, we investigate some problems concerning identity management for e-services and outline the next-generation identity management systems comparing them with today's Public Key Infrastructure (PKI) solutions.

Introduction

The widespread diffusion of on-line services provided by public and private organizations, firstly driven by e-commerce and more recently by e-government applications, has stressed the need of secure ways to *authenticate* users who need to access on-line resources. Kerberos [14], proposed at the beginning of the '90s, is a well-known example of a successful technology for authenticating users requiring access to resources belonging to a single organization. Later, Web-based systems created application environments that cross the boundaries of real organizations, posing new interoperability and scalability challenges, among which single-sign-on, and credential-based authentication plays an important role. This raised the need for cross-domain *digital identities*, as well as for standard procedures for user authentication to be adopted when on-line services, resource stake-holders, and users are geographically and organizationally distributed. A first answer to this new requirement was the Public Key Infrastructure (PKI) [3], today a well-known method for providing credential-based authentication and digital signatures solutions to electronic business and government applications. Specifically, a PKI is a standard set of technologies aimed at two main targets:

- definition of *hierarchies of certification authorities* (CAs) - mapping hierarchical administrative relationships, such as the one between a corporate head quarter and its corporate branches;

- support of *cross-certificates* mapping peer-to-peer relationships between cooperating parties, such as the one between partners of a joint consortium.

When early e-government systems were being designed, PKIs appeared to be well suited to satisfy e-gov requirements, such as the possibility for national or regional governments to establish root

[1] University of Milan, 26013 Crema – Italy, ardagna@dti.unimi.it
[2] University of Milan, 26013 Crema – Italy, cremonini@dti.unimi.it
[3] University of Milan, 26013 Crema – Italy, damiani@dti.unimi.it
[4] University of Milan, 26013 Crema – Italy, decapita@dti.unimi.it
[5] University of Milan, 26013 Crema – Italy, samarati@dti.unimi.it

CAs aimed at improving the interoperability among public agencies, guaranteeing correct procedures, and providing users of public services with a generalized token (e.g., identity certificate, smart card) that can be used as a credential for all possible requests. Also, PKIs usage was soon regulated by a number of international institutions.

Today, however, traditional PKIs look too complex as infrastructures for most e-government applications and several drawbacks have been identified. Many researchers share the view that a more flexible and cost-effective solution could be achieved following the *digital identity management* (DIM) approach. For the purposes of this paper, the term digital identity will be used to refer to two (non-disjoint) concepts: *nyms* and *partial identities* [7]. Nyms can be used to give a user a different identity under which operates at any interaction. A partial identity is any subset of the properties (e.g., name, age, credit-card, employment, and so on) associated with a user. Recently, some identity management solutions have been proposed such as the the Liberty Alliance's Identity Federation Framework (ID-FF), an open architecture and a set of specifications to enable federated identity management (`www.projectliberty.it`), the Oracle Identity Management (`www.oracle.com`), and the Microsoft .NET Passport (`www.passport.com`).

In the remainder of this paper we compare the public key infrastructures and the concept of digital identity and describe our current effort within the PRIME project.

PKI and identity management

In today's e-government systems, conventional PKI's basic privacy and authentication techniques [12] have been straightforwardly applied to bilateral communications, such as a citizen paying a local tax or a fine to a local administration. PKI authentication has also been used for multilateral communications, when a citizen signs a document or an application that several people will read. In many e-government applications, PKI is also used as a way to enhance privacy, by having the citizen encrypting the information she is submitting to the remote service with the recipient's public key. However, this approach is only suited to two-way communications. In the multilateral scenario, where multiple recipients share a message that should be kept private, current PKI technology does not provide a simple answer. Another drawback is the possibility for the citizen of losing the support holding her private key, required for decrypting information.[6] Moreover, if a private key is compromised, PKI certificates must be revoked and new ones issued, along with a new private key. Normally, private keys are split into several pieces, called *shares*, stored in different trusted locations. E-government applications relying on PKI on a large scale and over a significant period of time need advanced capabilities of managing end-users' keys lifecycle, including share management and provisions for recovering lost keys.

Privacy issues

Besides the drawbacks outlined above, some major privacy-related concerns have been raised about PKI as it does not provide a comprehensive solution for avoiding unauthorized disclosure of personal information. Indeed, personal credentials provided by users to service providers should be used for the sole purpose of granting access to the specific on-line service they are submitted to. Instead, they have been often used to profile users for marketing campaigns. To keep unauthorized disclosure of personal information in check, besides the adoption of specific legislation, the notion

---

[6] This problem is even worse with PKI than with traditional symmetric encryption, because the citizen is the only one who has access to her private key.

of digital identity itself is evolving beyond PKI. An environment for managing digital identities should support these 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.

- *Anonymity support*. 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.

With respect to these requirements, the usual way of designing PKI-based authentication and authorization systems is not satisfactory. In particular, selective disclosure of credentials is normally not implemented, because users attributes, either inserted into the X.509 identity certificate or collected as attribute certificates [10], 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. Pseudonymity, multiple identities and anonymity are also usually not supported in PKI-based architectures. Also, extensibility of the X.509 certificate format has encouraged the practice of encapsulating information needed for authorization within the identity certificate, making it difficult to cleanly separate the two sets of information. Furthermore, even when identity certificates and attribute certificates are disjoint, there has been a trend towards designing authorization architectures that strictly integrate the two types of certificates by referencing identity certificates inside attribute certificates.

These new requirements regarding digital identities have driven a number of new research projects . In the following, we describe the preliminary results of our ongoing activity in the framework of the PRIME project [16] (*Privacy and Identity Management for Europe*), funded by the European Commission. The PRIME project is a large-scale research effort aimed at developing an identity management system satisfying the requirements expressed for protecting users personal information and providing, at the same time, a framework that can be smoothly integrated with current architectures and on-line services.

A new vision of privacy and digital identity

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

- *Anonymity and end-user control*. The access control system should enable full end-user control over digital identity to be used. In other words, access control needs to operate even when interacting parties wish to remain anonymous or to disclose only specific attributes about themselves.

- *Flexible and expressive access control rules*. The access control rules should be able to refer to the different 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-government applications [6].

- *Client-side restrictions*. In addition to traditional server-side access control rules, users should be able to specify restrictions on how the released information can be used by their remote counterpart.

To take these issues into account, a new privacy-aware access control model is needed together with an access control protocol for the communication of policies and of identity information among parties. Specifically, we have introduced the definition of four different types of privacy policies.

- *Access control policies* govern access/release of data/services managed by the party (as in traditional access control).

- *Release policies* govern release of properties, credentials, *personal identifiable information* (PII) of the party and specify under which conditions they can be disclosed.

- *Sanitized policies* provide filtering functionalities on the response to be returned to the counterpart to avoid release of sensitive information related to the policy itself.

- *Data handling policies* define the personal information release will be (or should be) deals with at the receiving party.

In the next sections, we shall focus on access control policies, outlining their structure and underlying model.

A privacy-aware access control policy

Although the specific syntax of the access control rules will depend on the language used to define a policy, the policy has to contain the following basic elements.

- *Subject expression*. To provide expressive power and flexibility, a rule should specify the entities against which access must be controlled through expressions. Each expression identifies a set of subjects having specific properties. Each user is then associated with a *profile* that defines the name and value of some properties that characterize the user.

- *Object expression*. The characterization of the entities to be protected should be specified through expressions. As for subjects, each object is associated with a profile which defines the name and value of some properties that characterize the object.

- *Actions.* Policies must be able to make distinctions based on the type of action being performed (e.g., read, write, execute, and so on).

- *Purposes.* Data access requests are made for a specific purpose or purposes, which represent how the data is going to be used by the recipient.

- *Conditions.* Rules can include additional conditions, much in the same way as legislation often makes statements based on specified conditions.

- *Obligations*. To improve privacy, users can define some obligations attached to the data. Therefore, when a certain access is allowed, the parties involved must take some additional steps, following the defined obligations.

Each access request submitted to the system results in an *access decision* notifying that the request is *granted*, *denied*, or *undefined*. In particular, an undefined access decision is returned when the current information is insufficient to determine whether the request can be granted or denied and additional information is needed. As an example, suppose that a user can access a service if she is at least eighteen and can provide a credit card number. Three cases can occur: *i)* the system knows all the requested information and returns a positive response; *ii)* the system knows that the user is not yet eighteen and therefore returns a negative response; *iii)* the user has, for example, proved that she is eighteen and the system returns an undefined response together with the request to provide the number of a credit card.

An early prototype

We are now ready to describe the proof-of-concept prototype we developed to show how access control can work with minimal information disclosure on the part of the client. Our prototype supports only a subset of the requirements illustrated in the previous section. In particular, our current prototype deals with resource protection only and does not yet take into account obligation and purpose, which will be added in future releases. Compared to early commercial implementations of identity management systems like the one being developed at Lawrence Livermore National Laboratory [15], our approach provides the general notion of a policy language and of controlled release of personal information. On the other hand, at this stage we did not address engineering issues regarding scalability and backward compatibility with current standards.

Upon the submission of an access request from a client (possibly anonymous or non authenticated), which is characterized by the *subject* making the request, the *action* being requested, and the *object* (resource) on which the subject wishes to perform the action, the system returns the conditions (properties/credentials) that, according to the specified policies, the client should satisfy to gain the access (partial policy evaluation). Access control rules supported by the prototype have the following form:

subject `WITH` *subject expression* `CAN` *action* `ON` *object* `WITH` *object expression* `IF` *conditions*

where: *subject* identifies the subject to which the rule refers, *subject expression* is an expression that allows the reference to a set of subjects depending on whether they satisfy given conditions, *action* is the action to which the rule refers, *object* identifies the object to which the rule refers, *object expression* is an expression that allows the reference to a set of objects depending on whether they satisfy given conditions and *conditions* is a boolean expression of conditions that an access request, to which the rule applies, has to satisfy. Our prototype recognizes users specified in a *i) subject ontology* and for which there is a *ii) profile*. The subject ontology contains terms that can be used to make generic assertions on subjects. It recognizes also objects specified in an *i) object ontology* and for which there is a *ii) profile*.
The object ontology contains domain-specific terms that are used to describe the resource content and to make generic assertions on objects. The field *conditions* allows only credential-based conditions based on a credential ontology that defines abstractions and how these abstractions are implied by a combination of different credential types. A credential ontology is a set of facts of the form *abstraction* `IMPLIEDBY` *expression*, where *expression* can be a boolean formula of abstractions and/or credential types. For instance, a credential ontology can include the fact

`photo_id IMPLIEDBY driver-license` OR `passport` OR `id-card`. We assume that credential-based conditions are intended to be evaluated against the subject's profile.

The prototype includes four Prolog modules and a Java application interface. One Prolog module contains the evaluation engine, the other three modules contain the declarations of the ontologies, policies, and profiles.

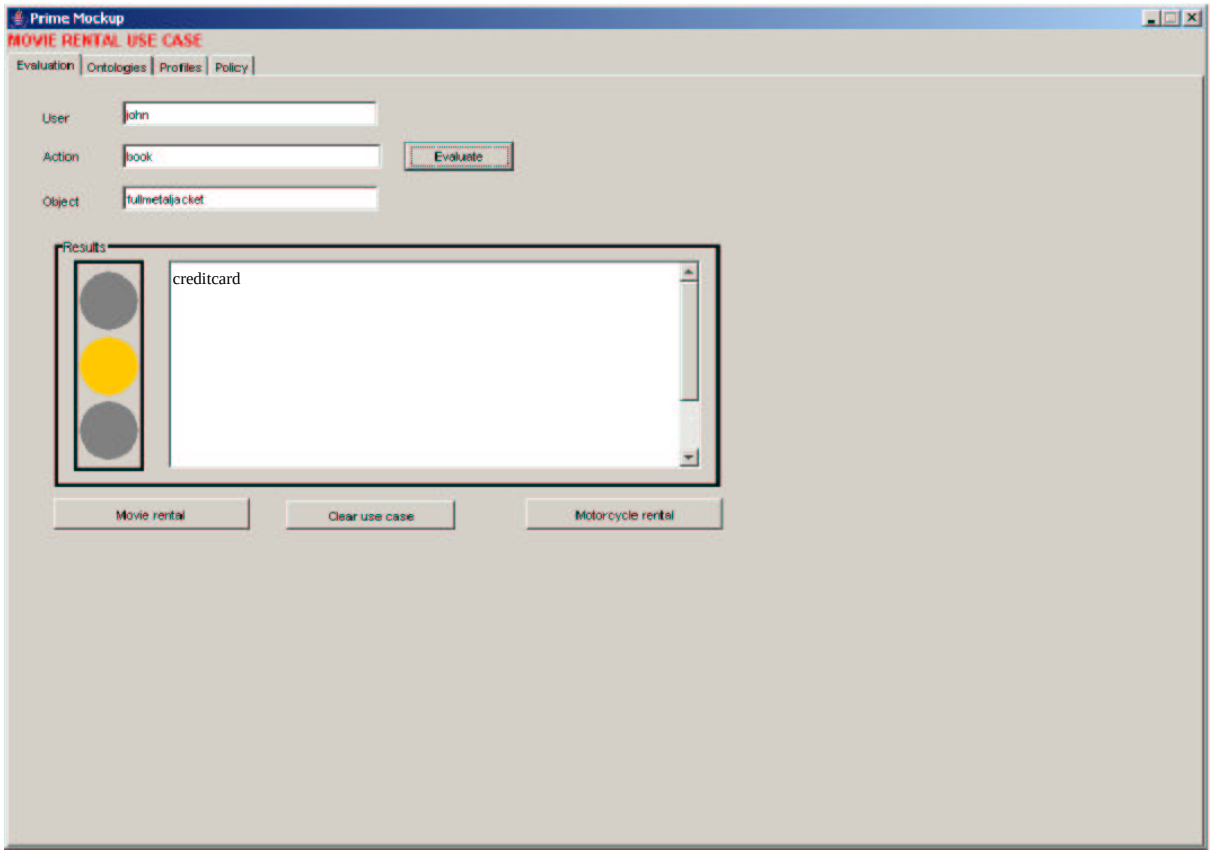

Figure 1 - Java Application Interface

The Java application interface (see Figure 1) has three text boxes, labeled `User`, `Action`, and `Object`, used to insert an access request. By clicking the button labeled `Evaluate`, the access request is evaluated against the defined access control policies. The output of the evaluation is shown in the box labeled `Results`, where a `3 color traffic light` indicates whether the request is denied (*red*), granted (*green*), or undefined (*yellow*). In the latter case the text area shows a list of alternatives (conditions) that must be fulfilled to gain the access.

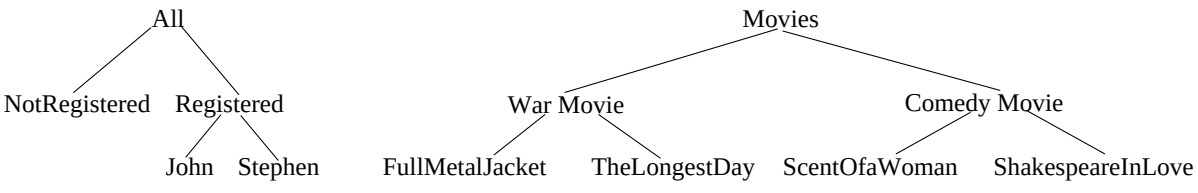

Subject Ontology **Object ontology**

Figure 2 - An example of subject ontology and object

As an example, consider a movie rental scenario that simulates the case of a user that wants to book a movie: Figure 2 illustrates an example of subject ontology and object ontology for this scenario.[7] Suppose now that we have defined the following access control rules:

Anonymous users can book movies if they provide a credit card.
rule 1: `anonymous WITH nocondition CAN book ON Movies WITH nocondition IF CreditCard`

War movies available on a web interface can be booked by registered users who live in Italy if they provide a credit card.
rule 2: `RegisteredUsers WITH nationality=Italian CAN book-online ON War Movies WITH availability=on-line IF CreditCard`

User `John` wants to `book` movie `fullmetaljacket` and we only know that `John` is Italian (from his user profile). Rule 1 is not applicable. The evaluation of condition `CreditCard` in rule 2 against the profile associated with user `John` is undefined: the system returns to `John` the information that is necessary to take a yes or no decision (see Figure 1).

Conclusions and Future Work

The protection of privacy in today's global infrastructure requires the combined application solution from technology (technical measures), legislation (law and public policy), and organizational and individual policies and practices. This paper has illustrated the main characteristics of the next-generation identity management systems and has presented the preliminary results of our ongoing activity in the framework of the PRIME project. Issues to be investigated in the access control area include: addition of a negotiation process among parties on the policies that should be applied on the collected data; extension of the notion of subject ontology to capture more complex assertions on subjects as well as the notion of object ontology and credential ontology; addition of obligations and purposes.

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.

Bibliography

[1] C. A. Ardagna and S. De Capitani di Vimercati. A comparison of modeling strategies in defining xml-based access control languages. CSSE, 2004.

[2] C. A. Ardagna, E. Damiani, S. De Capitani di Vimercati and P. Samarati. A Web Service Architecture for Enforcing Access Control Policies. In *workshop VODCA inside the International School On Foundations Of Security Analysis And Design*, 2004.

[3] A. Arsenault and S. Turner. Internet X.509 Public Key Infrastructure: Roadmap. Internet Draft, *Internet Engineering Task Force*, 2002.

[7] For simplicity, our prototype supports only assertions of the form *subject1* (*object1*, resp.) ISA *subject2* (*object2*, resp.) corresponding to the traditional abstractions that can be defined within the user and object domains, respectively.

[4] M. Blaze, M. Feigenbaum, J. Ioannidis and A. Keromytis. The KeyNote Trust Management System, version 2. Request For Comments 2704, *Internet Engineering Task Force*, 1999.

[5] 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.

[6] E. Damiani, A. Corallo, G. Elia. A Knowledge Management System Enabling Regional Innovation. Proc. of the VI international conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2002), Sep 16-18 2002, Crema (Italy).

[7] E. Damiani, S. De Capitani di Vimercati and P. Samarati. Managing Multiple and Dependable Identities.\newblock *IEEE Internet Computing*, November-December 2003.

[8] 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.

[9] C. Ellison. SPKI Requirements. Request For Comments 2692, *Internet Engineering Task Force*, 1999.

[10] S. Farrell and R. Housley. An Internet Attribute Certificate for Authorization. Request For Comments 3281, *Internet Engineering Task Force*, 2002.

[11] 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.

[12] R. Housley, W. Polk, W. Ford and D. Solo. Internet X.509 Public Key Infrastructure: Certificate and Certificate Revocation List (CRL) Profile. Request For Comments 3280, *Internet Engineering Task Force*, 2002.

[13] ITU Telecommunication Standardization Sector (ITU-T). Information Technology - Open Systems Interconnection – The Directory: Authentication Framework. Recommendation X.509 (03/00), *International Telecommunication Union*, 2000.

[14] J. Kohl and B. C. Neuman. The Kerberos network authentication service (version 5). Request For Comments 1510, *Internet Engineering Task Force*, 1993.

[15] M. Miley. Know Who, Know How: Using Oracle Technologies to Manage Identities. Oracle Magazine, n.8, 2004.

[16] PRIME (Privacy and Identity Management for Europe), http://www.prime-project.eu.org.

[17] 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.