An Infrastructural Approach to Secure Interoperability of Electronic IDs: the Bridging Backbone

Franco Arcieri[1], Andrea Dimitri[1], Fabio Fioravanti[1,2], Enrico Nardelli[1], Katia Pallucca[1], Alberto Postiglione[1,3] and Maurizio Talamo[1]

[1] NESTOR - Multimedia services security and certification Laboratory, University of Rome "Tor Vergata", Italy. fioravanti@nestor.uniroma2.it
[2] Dept. of Sciences, University of Chieti-Pescara, Italy
[3] University of Salerno, Italy.

Abstract. In this paper we propose a solution to secure interoperability between electronic ID management infrastructures enabling the provision of cross-border eServices to mobile citizens. Our proposal considers an interoperability architecture based on a *federation of national infrastructures* and follows a *cooperation based approach* which is fully compatible with and respectful of organizational and technical independence of existing national systems.

Provision of cross-country services to mobile citizens requires involved national citizen eAuthentication infrastructures (CEIs) to cooperate for establishing the identity of citizens, without interfering with organizational and technical solutions adopted in each CEI. Indeed, for obvious organizational and political reasons, there cannot be a single organization which is responsible for securing the authentication process in its entirety: instead, national organizations must keep full autonomy and responsibility for authenticating accessing citizens. Interoperability cannot be based on imposing a common technical architecture.

In order to be effective, technical solutions must satisfy this highly critical organizational constraint. As a consequence, any strongly centralized solution for addressing secure interoperability issues, although technically feasible and usable in strictly hierarchical environments (like, for example, some multinational companies), is not satisfactory in this more general context and is doomed to fail.

Moreover, different national organizations usually adopt different organizational schemes: for example, in some states there exists a single organization which is responsible for issuing national eIDs and validating foreign eIDs (centralized management of validation), while in other states these duties are distributed among different autonomous organizations, even from different public administrations (distributed management of validation).

These requirements strengthen our argument in favor of a *cooperation based approach* and an interoperability architecture which is based on a *federation of national infrastructures*, where the rules governing the federation dictates roles and responsibilities of each involved member State and conditions for their delegation and relying.

The cooperative approach is also followed by similar e-government projects which are currently being deployed in the United States [13,8,7].

Our solution to secure interoperability, by delegating the responsibility for citizen authentication to national CEIs, also satisfies the crucial requirement of preserving full compatibility with legacy eAuthentication systems.

The paper is organized as follows. In Section 1 we present our reference architecture for interoperability; in Section 2 we describe interoperability scenarios in eService provision to mobile citizens. Our solution to secure interoperability, the bridging backbone, is illustrated in Section 3, and in Section 4 we describe some peculiar functionalities it provides.

1 Reference Architecture

The architectural model adopted in our approach is inspired by the CEN's eAuth CWA [1] and considers the following three layers where interoperability has to be tackled: the *citizen device* layer, the *infrastructure* layer and the *application* layer.

The citizen device layer is the physical environment where the device is operating while accessing the infrastructure. The device can either be a smart-card (as it is common in many member States) or any other device supporting strong authentication (like mobile phones with cryptographic capabilities). Accessing devices must be linkable to personal identity, if needed, but should also be detachable from it in cases where only role identification is performed or where a certain degree of anonymity has to be guaranteed.

The infrastructure layer includes every component ranging from the physical interface with the citizen device to communication networks and systems, up to remote servers, including (i) a *user access point*, that is the local (w.r.t. the citizen device) part of the infrastructure, used by the citizen device for accessing the system, (ii) an *eService access point*, that is the remote part of the infrastructure, where the service providers system components interface with the infrastructure, and (iii) *validation services*, supporting eAuthentication procedures.

The application layer contains the applications which deliver services to users accessing the system by using citizen devices. For our purposes we will focus on applications requiring user authentication and, possibly, authorization credentials.

Our approach to interoperability considers three levels of functionalities, with increasing complexity, where interoperability has to be provided:

- identification and authentication: that is the process of associating a personal identifier with a citizen (identification) and proving trustworthiness of such association (authentication);
- authorization: that is the process of deciding whether to permit a particular action based on an identifier;
- electronic signature: that is the process of establishing authenticity of data and identity of the signer mainly for the purpose of producing verifiable records of transactions.

True interoperability can only be achieved if there exists a clear mapping between trust levels in various CEIs. This involves the definition of various degrees of trust existing in each CEI so that each device/system/service, once correctly placed with respect to this classification, can interoperate with other devices/systems/services with a correct understanding of mutual levels of trust during the interaction. For example, an authentication mechanism based on username and password issued during a completely online process cannot share the same level of trust of an authentication method based on credentials issued after careful physical identification.

Another important issue to be solved, which is orthogonal to interoperability of CEIs, is interoperability of electronic signatures. This is a difficult and very controversial point: for example, recommendations for management of electronic signatures contained in a EU directive have been interpreted in very different ways by different member states, and its use in some states as a basis for performing eAuthentication will surely be a source of potentially never-ending legal and juridical disputes. Furthermore, interoperability of security policies, user profiles and certificate validation are critical elements.

Future adoption of privilege management systems for handling authorizations and access rights of citizens depending on their role (doctor, policeman, CEO) within an organization, rather than on their identity alone, will also raise non-trivial semantic interoperability problems. Although this is an important research area to be investigated, role-based privilege management systems can only be built on top of effective and reliable authentication systems, and must be kept separate from them.

2 Interoperability Scenarios and Problems

Different interoperability scenarios can be envisaged, depending on whether the citizen device layer, the infrastructure layer and the application layer, are *on-us* (meaning national/domestic) or *not-on-us* (meaning foreign/alien). It is easy to imagine that different scenarios give rise to technical and organizational interoperability problems of different dimensions and nature.

Once a value is fixed for one layer, all possible interoperability scenarios are clearly identified. For example, the Italian access network (which means that the citizen device is physically in Italy) has to provide access in the five interoperability scenarios listed below:

- Italian devices accessing Italian services,
- Italian devices accessing foreign services,
- foreign devices accessing Italian services, and
- foreign devices accessing foreign services. This has two sub-cases:
  - foreign devices accessing their national services, and
  - foreign devices accessing services provided by a different foreign country.

As a further example of how this approach to modeling interoperability scenarios works, here below you can see a table showing the possible kinds of interoperability scenarios obtained by choosing the citizen device belonging to the Italian domain.

<table>
<tr><th></th><th>Italian application</th><th>foreign application</th></tr>
<tr><td>Italian infrastructure</td><td>Italian CEIs</td><td>partial interoperability</td></tr>
<tr><td rowspan="4">foreign infrastructure</td><td rowspan="4">partial interoperability</td><td>same CEIs</td></tr>
<tr><td>partial interoperability</td></tr>
<tr><td>different CEIs</td></tr>
<tr><td>full interoperability</td></tr>
</table>

A synthetic characterization of the different kinds of interoperability scenarios follows:

- *national CEIs*: components in all three layers belong to the same domain,
- *partial CEIs interoperability*: components in two layers belong to the same domain,
- *full CEIs interoperability*: components in each layer belong to different domains.

From an architectural viewpoint it is also important to identify interfaces lying between different components. Indeed, in the case of interoperability of national CEIs, some of them are managed by different national systems, raising the need for establishing a secure and reliable dialog among them.

A first interface is between the citizen device requesting access to the system and the physical devices communicating with it at the user access point.

A second interface is between the user access point and the service access point, that is between the local terminal application and the access point to the requested service.

A third interface is between a user or service access point and the validation service used to verify the validity of credentials presented by users. This interface is highly critical for services requiring user authentication.

The fourth interface is between a service access point and an eService.

An example scenario at the widest possible interoperability level is the following: an Estonian citizen wishes to make access to an Italian service while visiting Belgium. In this case, the first interface is physically in Belgium, the second interface spans the three countries, the third interface is physically in Estonia, and the fourth is in Italy.

Problems of high technical complexity stem from the need of managing the whole process in an efficient and effective manner, while ensuring, at the same time, interoperability of geographically distributed IT-based systems, independently of technical solutions used by participating organizations, and fulfillment of privacy and security constraints in a democratic manner [10,9]. There are in fact two critical functional capabilities for the interoperability architecture: security and performance.

The first delicate functionality is end-to-end security. By this term it is meant the capability of ensuring traditional security requirements (from basic ones: confidentiality, integrity, authentication, authorization, to derived ones: auditing, non repudiation, etc.) from the citizen accessing devices all the way down to the point providing the required service.

The second one is performance experienced by end-users. Due to the cooperative approach that has to be followed in designing the overall interoperability architecture and to the size of federated systems of national CEIs that will result, each service invocation may require establishing and traversing several times geographically long and organizationally complex communication paths. A carefully designed architecture must be able to cache information at usage points and keep it fresh to avoid the well-know attacks based on exploiting stale security information.

3 Our Solution: the Bridging Backbone

Efficient satisfaction of the end-to-end security requirement requires the definition of a highly secure and efficient exchange layer among national CEIs, allowing Management Centers of the involved CEIs to quickly exchange all information required to properly authenticate accessing citizens and enabling an efficient and secure management of service provision. This exchange layer is overlaid to and logically distinct from existing CEIs. It is important to stress that, for reasons of efficiency, only cross-border interactions which are important for security and privacy purposes or for documenting interaction between CEIs, will be required to pass through this overlaid layer.

During provision of cross-border services to mobile citizens, national systems used for managing electronic IDs interoperate at different layers as follows: after citizen credentials are obtained from the citizen device and understood by the access network, the citizen's request is first relayed to the competent national infrastructure Management Center for authentication and afterwards to the management centers which are competent for providing the requested services. National CEIs interface each other through a highly secure communication network which is logically distinct from national access networks. Our proposal to techno-organizational problems in secure interoperability among CEIs will thus be based on defining a permanent infrastructure layer, called *bridging backbone*, providing security services to interactions between Management Centers of national CEIs. Each of these Centers will continue to operate in the normal way under the "national CEIs" scenario, while it will cooperate with the other Centers under the other two interoperability scenarios, each working within its responsibility boundaries.

In our approach, security functions have to be based on a permanent infrastructure layer, since this is the only approach which is able to guarantee, at a reasonable cost, efficiency of eService provision and effectiveness of security in open and intrinsically insecure environments like the Internet. In other words, we do not deal with security functions within application, but consider them as infrastructural services, much in the same way communication services are nowadays considered: from the application viewpoint, in fact, details regarding how messages are transported by the communication network to their destination are completely transparent. In the same way, applications in our architecture do not take care of the management of security functions, which are instead provided by an independent layer put on top of the layer providing communication services.

A mandatory requirement of the cooperation based approach is the ability to document transactions that were carried out during interaction between national Management Centers. Given the legal value attached to data being managed and exchanged in this process and the fact that many various kinds of mistakes can take place during the interaction, it is necessary to clearly and unequivocally understand who did what. The absence of a super-national organization that can supervise and direct the activity of national Management Centers makes these certification functions a mandatory requirement. Moreover, as certification functions in a federation of national infrastructure play a back-office and subordinate role, they are fully acceptable by involved organizations, both from political and organizational viewpoints.

It is important to stress that in the real world of non-electronic services and whenever some kind of contractual responsibility is involved, security functions are always based, to various degree, on some form of permanent infrastructure. For example, public utilities like power supply, water, and sewage are provided by Municipalities to houses on the basis of the house ownership or renting. People interact with banks in buildings and offices clearly and permanently identifiable as bank settings (even ATMs are usually placed in trustable environments). Also e-banking, the currently most widespread eService among the ones where trust is a fundamental aspect, is based on an initial setup phase where a security infrastructure is established: the customer goes physically to branch offices for signing the contract and receives codes and instructions for accessing the service on the Internet.

A further important point regarding security in interaction between institutions (as compared to interaction among people) is that organizations typically do not allow any inside member to unilaterally establish trust with external entities. The reality of institutional cooperation shows that inter-institutional trust is always based on bilateral agreement at the organizational level. The electronic counterpart of this convention is that, at the IT level, there must be an infrastructure layer providing security functions, and security functions are provided with reference to and after that an agreement is formally in place between the involved organizations.

This approach gives maximum flexibility to each involved organization, by respecting its techno- organizational choices, and allows to design and build a scalable and efficient system, because interoperability is not based on country-to-country system interfaces deriving from bilateral agreements. The bilateral agreements approach would be, in fact, viable and effective only when there are very few actors: as soon as the stakeholders are more than three or four its complexity becomes unmanageable.

4 Management of Security/Authorization Services on the Bridging Backbone

The management of authorization rights is a further critical element for interoperability. Traditionally, the task of authorizing access requests to a given service is accomplished by the service provider itself, which ultimately holds the responsibility for service provision, but, in geographically dispersed scenarios, performance considerations may suggest a pushing solution where preliminary authorization information is properly and progressively moved towards the front-end side of the architecture (i.e., the system access points). Critical issues of efficient distribution and update of authorization data derive, that can be solved by using approaches similar to those adopted for improving efficiency of PKI certificate validation.

In the following we consider an authorization model, inspired by the requirements which led to the development of languages and models [12,11,14,2] for management of authorization rights in distributed network environments.

In the considered authorization model we identify three main entities: *subjects*, *resources* and *actions*. Entities are specified by a set of attributes of the form $\langle n, v \rangle$, where n is the attribute name and v is the attribute value. Attributes values can be statically defined or can dynamically change over time.

A subject is an entity which wants to perform an action on a resource. Subjects can be either accessing citizens or processes acting on citizens' behalf.

A resource is a computational entity which is available for use to authenticated subjects on the interoperability architecture. A resource is typically specified by using the following information: the name of the resource, the host where it is located, and the application protocol and the network port which must be used for accessing the resource. Example of resources are FTP directories, file systems and eService applications.

An action is an operation which a subject wants to perform on a resource. Actions can be specific to a particular resource or application protocol and thus, not all actions can be performed on a given resource. The complete set of actions which can potentially be performed on a resource must be explicitly agreed upon by the organizations involved and stated in formal agreement documents. Example of actions are the following: read, write, execute, HTTP GET, HTTP POST.

Authorization policies are sets of authorization rules which specify if and how subjects can perform actions on resources by using constraints on resource and action attributes (f.e. "only HTML documents can be accessed"), or context constraints (f.e. "access is granted between 9 AM and 7 PM", "access is granted only twice a day", "access is limited to a maximum of two concurrent users"). In order to simplify the definition of authorization policies, authorization rules need not refer to every particular instance of subjects, resources or actions but can refer to classes of subjects. When evaluating an authorization policy, rules may conflict and can be combined in different ways. We assume that the authorization rights granted by an authorization policy consists of the union of the authorization rights granted by each applicable authorization rule therein contained.

We now illustrate a scenario where authorization rights are statically defined and cannot be changed, and a scenario where authorization rights can be dynamically modified. In both cases, users are authenticated by the bridging backbone at the beginning of the session, by performing a single sign-on procedure. Notice that, by following this approach, each organization still retains full control of the hosts operating on the network, and, as already mentioned, no changes to applications or to intra-organizational architectures are required.

In a static scenario, the set of authorization rights owned by users is statically determined by the configuration of the bridging backbone, and does not change during time unless the bridging backbone is externally reconfigured by manual intervention. The set of authorization rights owned by a user consists of direct authorization rights, which are explicitly granted by the bridging backbone, and derived ones, which are obtained by delegation and relying.

We can model authorizations by using a labeled directed graph whose nodes are labeled by subjects and resources and whose edges are labeled by actions. The set of authorization rights owned by each subject can be represented as the transitive closure of the peer-to-peer authorization relation starting from authorizations directly owned by the subject: a subject s can perform an action a on a resource r iff there exists a path from s to r labeled $a_0, \ldots, a_n$ where, for all $i = 1, \ldots, n$, if action a_i is performed then action a_{i+1} can also be performed.

In the initial state of the dynamic scenario, i.e. soon after the bridging backbone is set up, the only interactions allowed by the bridging backbone are authorization requests from accessing citizens to the Policy Decision Point, which evaluates authorization policies and replies with authorization decisions.

When a subject wants to perform an action on a resource, it sends an authorization request containing attributes of the resource, action, and other related information to the authorization server. The authorization server, upon receiving the request for authorization, examines it and retrieves policies which are needed for determining whether and how to satisfy the request. As a result of the decision process, the authorization server sends back to the requesting subject a response containing the authorization decision. If the authorization request is accepted the authorization server proceeds in activating a procedure which reconfigures some software components affecting the behavior of the bridging backbone. Only after this reconfiguration process is successfully completed, the subject is allowed to establish a secure communication channel with the resource over the bridging backbone and to perform operations which are compliant with authorization policies agreed upon by involved organizations. The secure channel is destroyed at the end of the work session or after a timeout occurs.

5 Conclusions

In this paper we propose a solution to security and interoperability issues raised by management of electronic IDs in provision of cross-border eServices to mobile citizens: the bridging backbone. Our proposal considers an interoperability architecture based on a *federation of national infrastructures* and follows a *cooperation based approach* which is fully compatible with and respectful of organizational and technical choices of existing systems. Indeed, the bridging backbone provides infrastructural security services (like confidentiality and integrity services, authentication, authorization and auditing), in an easy and transparent manner, independently from locally deployed network technology and topology. The bridging backbone provides security services as a layer lying between the application and communication layers, which is in charge of monitoring network connections and securing them according to the cooperation policies of the federation of involved infrastructures (see [5,4,6,3] for details).

In our view security is an infrastructural service of inter-organizational communication, not an add-on service. Our approach, by providing applications with security services in a completely transparent, infrastructural way, allows separating the issues related to security services and business logic, thereby reducing the risks of introducing security flaws. This is in contrast with the standard approach, where security services are usually provided at different levels of the protocol stack.

Moreover, we solve the organizational pitfalls of a naive use of PKIs, where trust can be established unilaterally, by allowing cooperation between members of different organizations only on top of the bridging backbone layer, which is set up only after a bilateral agreement is formally in place at the organizational level. Additionally, auditing services provided by the bridging backbone give the ability to certify successful e-services interaction and composition, to identify culprits of unsuccessful service provision, and to monitor actual performance of service provision [5].

References

1. CEN/ISSS Workshop on e-Authentication, 2004. http://www.cenorm.be/.
2. Gail-Joon Ahn. Specification and Classification of Role-based Authorization Policies. In *Twelfth International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises* June 09 - 11, 2003 Linz, Austria, 2003.
3. Franco Arcieri, Elettra Cappadozzi, Enrico Nardelli, and Maurizio Talamo. SIM: a Working Example of an E-government Service Infrastructure for Mountain Communities. In *Workshop Electronic Government (DEXA-eGov'01), associated to the 2001 Conference on Databases and Expert System Applications (DEXA'01)*, pages 407–411, Munich, Germany, September 2001. IEEE Computer Society Press.
4. Franco Arcieri, Mario Ciclosi, Fabio Fioravanti, Enrico Nardelli, and Maurizio Talamo. The Italian Electronic Identity Card: a short introduction. In *The National Conference on Digital Government Research (dg.o2004), May 24-26, 2004, Seattle, Washington, USA*.
5. Franco Arcieri, Fabio Fioravanti, Enrico Nardelli, and Maurizio Talamo. Inter-organizational E-Services Accounting Management. In *3rd IFIP conference on e-Commerce, e-Business, and e-Government (I3E-03)* Sao Paolo, Brasil. Kluwer Academic Publishers, September 2003.
6. Franco Arcieri, Fabio Fioravanti, Enrico Nardelli, and Maurizio Talamo. Reliable Peer-to-Peer Access for Italian Citizens to Digital Government Services on the Internet. In Roland Traunmüller, editor, *Electronic Government: Third International Conference, (EGOV'04),, Zaragoza, Spain, August 30 - September 3, 2004*, volume 3183 of *Lecture Notes in Computer Science*, pages 250–255. Springer-Verlag, 2004.
7. United States Federal PKI Operational Authority. Federal Public Key Infrastructure (FPKI) Architecture Technical Overview, 2005.
8. United States Federal PKI Policy Authority. X.509 Certificate Policy for the Federal Bridge Certification Authority, 2002.
9. Sokratis K. Katsikas, Stefanos Gritzalis, and Javier Lopez, editors. *Public Key Infrastructure: First European PKI Workshop: Research and Applications, EuroPKI 2004, Samos Island, Greece, June 25-26, 2004. Proceedings*, volume 3093 of *Lecture Notes in Computer Science*. Springer, 2004.
10. Sokratis K. Katsikas, Javier Lopez, and Günther Pernul, editors. *Trust and Privacy in Digital Business, First International Conference, TrustBus 2004, Zaragoza, Spain, August 30 - September 1, 2004, Proceedings*, volume 3184 of *Lecture Notes in Computer Science*. Springer, 2004.
11. OASIS. eXtensible Access Control Markup Language (XACML). http://www.oasis-open.org/committees/xacml/.
12. OASIS. Security Assertion Markup Language (SAML). http://www.oasis-open.org/.
13. United States General Accounting Office. Planned e-Authentication Gateway Faces Formidable Development Challenges, 2003. Report to the Committee on Government Reform and the Subcommittee on Technology, Information Policy, Intergovernmental Relations, and the Census, House of Representatives.
14. Ravi S. Sandhu, Edward J. Coyne, Hal L. Feinstein, and Charles E. Youman. Role-Based Access Control Models. *IEEE Computer*, 29(2):38–47, 1996.