Towards a Next-Generation Trust Management Infrastructure for Open Computing Systems

Position Paper

Yücel Karabulut

SAP Corporate Research
Vincenz-Priessnitz-Str.1, 76131 Karlsruhe, Germany.
yuecel.karabulut@sap.com

Abstract. Basically, there are two intertwined kinds of security mechanisms: monitoring including access control and cryptographic protocols. The purpose of an access control system is to enforce security policies by gating access to, and execution of, processes and services within a computing system. Specification and enforcement of permissions can be based on asymmetric cryptography. In order to employ asymmetric cryptography in open computing environments we need appropriate trust management infrastructures that enable entities to establish mutual trust. Management of trust is organized within a public key infrastructure, PKI for short. Credentials assert a binding between a principal, represented by a public key, and some property. Current proposals investigating the definition of PKI and the application of credential-based access control treat exisiting PKI models (e.g. X.509) and trust management approaches (e.g. SPKI/SDSI) as competing technologies. We take a different position. We argue here that a trust management infrastructure for open computing environments has to use and to link existing approaches. We explain which requirements a next-generation trust management approach has to fulfill. After presenting an application scenario, we finally outline the design of a next-generation trust management approach that we believe really would appear to be worthwhile for a broad spectrum of applications.

1 Introduction

The proper administration of computing systems requires to specify which clients are allowed to access which services, and to effectively and efficiently enforce such specifications. In a local computing system, a specification can be represented by traditional access rights granted to known identified individuals and thereby to the processes under their control. The enforcement is mostly based on identification and authentification of requesting individuals over a trusted physical path and on keeping track of the processes they are controlling.

In the Internet most interactions including business transactions occur between strangers, due to billions of spontaneous users and the fact that most

of them do not share a common security domain. Thus, Internet constitutes a global computing infrastructure in which entities need to reason about the trustworthiness of other entities in order to make autonomous security decisions.

In the modern computing environments [14] emerging from these trends, some basic assumptions of traditional access control approaches are not longer valid. Traditional access control mechanisms operate under a closed world assumption, in which all of the entities are registered and locally known. When the server and the client are unknown to one another and when resources are to be shared across administrative boundaries, the conventional authorization scheme fails. Thus, we cannot reasonably assume anything like a trusted physical path between remote agents.

In order to overcome these and related difficulties a diversity of proposals has arised. While all proposals exploit cryptography, some of them use symmetric cryptographic mechanisms, like Kerberos [13], and others rely on asymmetric cryptography, like X.509 [8, 9] and SPKI/SDSI [7, 6, 5]. Accordingly, we can specify and enforce permissions of clients on remote servers by employing modern access control approaches which are based on asymmetric cryptography. In order to employ asymmetric cryptography in open computing environments we need appropriate trust management infrastructures that enable entities to establish mutual trust. Management of trust is organized within a public key infrastructure, PKI for short. Credentials are digital and digitally signed documents that assert a binding between a principal, represented by a public key, and some property.

Current literature treat existing PKI models and trust management approaches as competing technologies even as dueling theologies [4]. We take a different position. We argue that a trust management infrastructure for an open and dynamic computing environment has to use and to link existing PKI models. Accordingly, we designed a hybrid PKI model to be used for specifying and enforcing permission in open computing systems. The hybrid PKI model, as reported in [2, 3], unifies and extends previous PKI approaches [9, 8, 6, 7, 5].

The sole purpose of this position paper is to stimulate discussion in a workshop on security and privacy in pervasive computing. In particular, it is not our goal here to put forth new results and proposals. All of the technical material alluded to here has been developed in previous work [10, 2, 12, 11].

2 Thoughts on a Next-Generation Trust Management Infrastructure

Requirements: As a basis for emerging distributed applications which aim to follow credential-based access control policies, we would like to see the following features[1] supported by a next-generation trust management infrastructure that enables interoperability between heterogeneous security domains:

[1] In order to concentrate on the main concepts, we didn't treat the topic of certificate revocation here. Instead we just assume appropriate mechanisms to handle this issue.

- free properties (e.g. personal data, a skill, group membership)
- bound properties (e.g. a ticket, a capability, a role)
- conversion of free properties into bound properties
- the model of trusted authorities with licencing (e.g. X.509)
- the model of owners with delegation (e.g. SPKI/SDSI)
- administrative properties (e.g. trustee, licensee, delegatee)
- recursive trust evaluation (e.g. path validation, chain reduction)
- expressive certificates or credentials
- expressive authorization policies based on role-based access control
- authorization decision engines
- credential management components

In addition to these features, the anonymity need of the clients has to be considered. While requesting accesses to the resources, clients may be unwilling to reveal their identities for private reasons and thus prefer to remain anonymous. Additionally for a resource owner, it may be necessary to see evidences of a client's eligibility rather than to know *who* they are. Thus, the trust management infrastructure should provide mechanisms (e.g. private credentials) to support anonymity of the clients.

An Application: A typical scenario exploiting the use of credentials for access control runs as follows. A client is represented by (one of) his public key(s) and characterized by the assigned properties. A resource owner follows a confidentiality policy that is expressed in terms of characterizing properties. An agent as resource owner receives a signed request together with a set of credentials stemming from the pertinent client. The agent firstly ensures the authenticity with respect to the bound public keys and with respect to the actual holder of the corresponding private key by applying appropriate challenge-response protocols and secondly evaluates his trust in the signing issuer. Then the agent decides on the permission of the request by evaluating the properties extracted from submitted credentials with respect to his confidentiality policy.

Depending on the application and the underlying trust relationships between the involved entities, such scenarios can be realized by employing different PKI models and trust management approaches. We see arguments of the style this-model-is-better-than-another-model. PKI trust relationships must be built on real-world[2] trust relationships. In many real-world scenarios, trust relationships consist of hierarchies, trust networks, and combinations of two. Therefore, we argue that a trust management infrastructure, as required by dynamic computing environments, has to use and to link both kinds of PKI models.

More concretely, we consider the following scenario. In [1], we proposed a secure information *integrating* mediation approach (*i*-mediation for short) considering the dynamics and conflicting interests of mediation participants. In mediated information systems, a client seeking information and various autonomous sources holding potentially useful data, are brought together by a third kind of

[2] It is also important to observe that in some cases, such as the use of PKI to support anonymity, it can be important to make sure that PKI trust relationships don't follow real world trust relationships.

independent components, called *mediators*. Data sources in i-mediation, following property-based security policies, aim at supporting a wide range of potential clients, which are in general unknown in advance and may belong to heterogeneous and autonomous security domains. This raises the challenge how remote and autonomous entities can agree on a common understanding of certified properties, and other issues related to these properties (e.g. encoding formats).

In such situations the sources wish to be assisted to determine potentially eligible clients. To reach potentially eligible clients, which might belong to remote security domains, the sources will need to *trust* mediating agents having the required domain expertise as well as the relationships with the potential clients. As a concrete solution, we proposed an additional mediation functionality, called entity *finding* mediation, f-mediation for short. F-mediation employs our hybrid PKI model [2].

Outline of the Infrastructure: In [2], we classified previous PKI approaches as based on trusted authorities with licencing and dealing with free properties (characterizing attributes including identities) and the corresponding certificates[3], e.g. X.509, or based on owners with delegation dealing with bound properties (including capabilities) and the corresponding credentials[4], e.g. SPKI/SDSI. We extended and integrated these approaches into a hybrid PKI model which uses protocols to convert *free properties* into *bound properties*. Furthermore, we unified licencing and delegation by introducing *administrative properties*.

An instance of the full hybrid PKI model consists of overlapping components of three kinds: a) trusted authorities (also called trustees) and licensees for and a holder of a free property together with a verifier of this free property, b) an owner and delegatees for and a grantee of a bound property, and c) a holder of free properties and a grantor of a bound property. The grantor follows a *property conversion policy* that maps free properties on bound properties, where the property conversion policy is a part of grantor's whole security policy. More precisely, the property conversion policy specifies *which set of free properties an entity has to enjoy in order to obtain a bound property assignment.*

A typical interaction for a property conversion process runs as follows: A holder of free properties requests a promise for a permission, i.e., a bound property. For this purpose, the holder shows her certified free properties and applies for a bound property from the grantor who is acting as an authorizer on behalf of and in explicit delegation of a resource owner. The grantor, after verifying the submitted free property-certificates with the supporting licences, applies his conversion policy on the free properties extracted from the submitted certificates, and finally, if all checks have been successfully completed, grants a bound property-credential where the subject (grantee) is the same as in the submitted free property-certificates.

Our hybrid PKI model brings together different PKI models and trust management approaches. The business advantage of such a model is clear. By em-

[3] X.509 uses the terms *public-key certificate* and *attribute certificate*.
[4] SPKI uses the term *authorization certificate*.

ploying a unifying PKI model, which provides a seamless interoperation between hetereogeneous and autonomous security domains, organizations can broaden their potential customer base and collaborators base.

3 Acknowledgements

It is a pleasure to thank Joachim Biskup with whom I've had extensive discussions about trust management, PKI models and secure mediation.

References

1. C. Altenschmidt, J. Biskup, U. Flegel, and Y. Karabulut. Secure mediation: Requirements, design and architecture. *Journal of Computer Security*, 11(3):365–398, 2003.
2. J. Biskup and Y. Karabulut. A hybrid PKI model: Application to secure mediation. In *16th Annual IFIP WG 11.3 Working Conference on Data and Application Security*, pages 171–182, Cambridge, England, July 2002. Kluwer Academic Press, 2003.
3. J. Biskup and Y. Karabulut. Mediating between strangers: A trust management based approach. In *2nd Annual PKI Research Workshop*, pages 80–95, Gaithersburg, USA, Apr. 2003.
4. B. Chinowsky. Summary of the panel discussion Dueling Theologies. In *1st Annual PKI Research Workshop*, Gaithersburg, Maryland, USA, Apr. 2002.
5. D. Clarke, J.-E. Elien, C. Ellison, M. Fredette, A. Morcos, and R. L. Rivest. Certificate chain discovery in SPKI/SDSI. *Journal of Computer Security*, 9(4):285–322, 2001.
6. C. Ellison. SPKI/SDSI certificates. `http://world.std.com/~cme/html/spki.html`, Aug. 2001.
7. C. M. Ellison, B. Frantz, B. Lampson, R. Rivest, B. M. Thomas, and T. Ylonen. Simple public key certification. Internet draft, work in progress. `http://www.ietf.org/ids.by.wg/spki.html`, June 1999.
8. IETF X.509 Working Group. Public-key infrastructure (X.509). `http://www.ietf.org/html.charters/pkix-charter.html`, 1998.
9. ITU-T. ITU-T recommendation X.509: The directory - public-key and attribute certificate frameworks, 2000.
10. Y. Karabulut. *Secure Mediation Between Strangers in Cyberspace*. PhD thesis, University of Dortmund, 2002.
11. Y. Karabulut. Developing a trust management based secure interoperable information system. In *Special Session on Security and Privacy in E-Commerce within the 6th International Conference on Electronic Commerce Research*, pages 342–359, Dallas, USA, Oct. 2003.
12. Y. Karabulut. Implementation of an agent-oriented trust management infrastructure based on a hybrid PKI model. In *1st International Conference on Trust Management*, LNCS 2692, pages 318–331, Heraklion, Crete, Greece, May 2003.
13. B. Neuman and T. Ts'o. Kerberos: An authentication service for computer networks. *IEEE Communications*, 32(9):33–38, Sept. 1994.
14. A. S. Tanenbaum and M. van Steen. *Distributed Systems: Principles and Paradigms*. Prentice Hall, Upper Saddle River, NJ, Sept. 2002.