CERIAS Tech Report 2004-78

HIDDEN ACCESS CONTROL POLICIES WITH HIDDEN CREDENTIALS

by K.Frikken, M. Atallah, J. Li

Center for Education and Research in
Information Assurance and Security,
Purdue University, West Lafayette, IN 47907-2086

Hidden Access Control Policies with Hidden Credentials *

Keith Frikken
Purdue University
1250 N. University Street
West Lafayette, IN 47907
kbf@cs.purdue.edu

Mikhail Atallah
Purdue University
1250 N. University Street
West Lafayette, IN 47907
mja@cs.purdue.edu

Jiangtao Li
Purdue University
1250 N. University Street
West Lafayette, IN 47907
jtli@cs.purdue.edu

ABSTRACT

In an open environment such as the Internet, the decision to collaborate with a stranger (e.g., by granting access to a resource) is often based on the characteristics (rather than the identity) of the requester, via digital credentials: Access is granted if Alice's credentials satisfy Bob's access policy. The literature contains many scenarios in which it is desirable to carry out such trust negotiations in a privacy-preserving manner, i.e., so as minimize the disclosure of credentials and/or of access policies. Elegant solutions were proposed for achieving various degrees of privacy-preservation through minimal disclosure. In this paper, we present an efficient protocol that protects both sensitive credentials and policies. That is, Alice gets the resource only if she satisfies Bob's policy, Bob does not learn anything about Alice's credentials (not even whether Alice got access or not), and Alice learns neither Bob's policy structure nor which credentials caused her to gain access.

Categories and Subject Descriptors

K.4.4 [**Computers and Society**]: Electronic Commerce—*Security*

General Terms

Design, Security

Keywords

Secure multi-party computation, access control, trust negotiation, hidden credentials, privacy

1. INTRODUCTION

Whereas in the past access decisions were based on the identity of the entity requesting a resource, in open systems such as the Internet, this approach becomes ineffective, as the resource owner and the requester usually belong to different security domains controlled by different authorities and are unknown to each other. The modern alternative is to use *digital credentials* for satisfying access policies. Digital credentials, the digital equivalent of paper credentials, are digitally signed assertions about the credential owner by a credential issuer. The decision to access a resource is based on the attributes in the requester's credentials, such as age, citizenship, employment, or credit status.

As a simple example where both the credentials and the policy are sensitive, consider an online business that grants access to media records by sending access keys to its client's special media-reader software – keys that the reader uses to "unlock" encrypted media records that are freely downloaded in encrypted form. Certain records are treated differently from the rest: The online business may grant access to these records only if the requester has a disability, or is a senior citizen, or is terminally ill, and has an income of under $30K a year. This requirement involves four attributes (denote them by a_1, a_2, a_3, a_4) and the policy is $(a_1 \vee a_2 \vee a_3) \wedge a_4$. In order to gain access to the sensitive records in Bob's database, Alice needs to show or prove to Bob that she satisfies the policy. However, neither Alice nor Bob is willing to disclose her/his private information. Alice does not want to reveal her credentials, as her credentials contain sensitive information about her (e.g., health, age, income, etc). Bob does not want to reveal the policy, even to those who satisfy the policy, so as to make it harder for an adversary to know which credentials he should forge or otherwise illicitly obtain.

In other examples, the motivation for hiding the policy is not security from an evil adversary, but simply the desire to prevent legitimate users from "gaming" the system – e.g., changing their behavior based on their knowledge of the policy (which usually renders an economically-motivated policy less effective). This is particularly important for policies that are not incentive-compatible in economic terms (they suffer from perverse incentives in that they reward the wrong kinds of behavior, such as free-loading). In yet other examples, the policy is simply a commercial secret – e.g., Bob has pioneered a novel way of doing business, and knowledge of the policy would compromise Bob's strategy and invite unwelcome imitators.

Finally, it is important to point out that a process that protects Alice's credentials from Bob is ultimately not only to Alice's advantage but also to Bob's: Bob no longer needs to worry about rogue insiders in his organization illicitly leaking (or selling) Alice's private information, and may even lower his liability insurance rates as a result of this. Privacy-preservation is a win-win proposition, one that is appealing even if both Alice and Bob are honest and trustworthy entities.

*Portions of this work were supported by Grants IIS-0325345, IIS-0219560, IIS-0312357, and IIS-0242421 from the National Science Foundation, Contract N00014-02-1-0364 from the Office of Naval Research, by sponsors of the Center for Education and Research in Information Assurance and Security, and by Purdue Discovery Park's e-enterprise Center.

WPES'04, October 28, 2004, Washington, DC, USA.

Recently, Holt et al. proposed a novel hidden credentials system [2] that protects sensitive credentials and policies. Hidden credentials are used in a way that they are never shown to anyone, thus the sensitive credentials are protected. The hidden credentials system also protects sensitive policies; however, we believe the protection is not enough for the following reasons. First, the policy structures are revealed in their system. Second, if an access to a resource is granted, the requester learns which attributes gave her access. Finally, even if the requester cannot access the resource, she might learn some partial information about the policy.

In this paper, we present an efficient protocol that protects both sensitive credentials and policies. That is, Alice gets the resource only if she satisfies the policy, Bob does not learn anything about Alice's credentials (not even whether Alice got access or not), and Alice learns neither Bob's policy structure nor which credentials caused her to gain access. Our protocol is built on the hidden credentials system [2] and scrambled circuit evaluation [3]. We next present our model and problem definition.

2. MODEL AND PROBLEM DEFINITION

We first briefly review the hidden credentials system [2], then define our problem. In the hidden credentials system, there is a trusted Credential Authority (CA) who issues credentials for users. Each user has a unique username and each credential contains a username and an attribute. Suppose the CA created an identity-based encryption (IBE) system [1], a hidden credential c for user nym and attribute $attr$ is the private key corresponding to $nym||attr$. We use $c.attr$ to denote the attribute correspoding to c.

The problem of protecting both sensitive credentials and policies is defined as follows. Suppose Alice has username nym and m hidden credentials $c_1, \ldots, c_m$; Bob has a resource M that Alice wants to access. The policy over M is a function $p(x_1, \ldots, x_n) : \{0,1\}^n \rightarrow \{0,1\}$ over n attributes $a_1, \ldots, a_n$. If Alice has attribute a_i in her credentials (i.e., there exists c_j, $1 \leq j \leq m$, such that $c_j.attr = a_i$), then x_i is set to be 1, to be 0 otherwise. The policy is satisfied if and only if $p(x_1, \ldots, x_n) = 1$. Our goal is that (1) Alice obtains M if her credentials satisfy Bob's policy, (2) Bob learns nothing about Alice's credentials, (3) Alice learns minimum information about Bob's policy.

3. OUR PROTOCOL

There are two primary phases in our protocol: (1) a credential hiding phase and (2) a blinded policy evaluation phase. During the credential hiding phase, Alice and Bob engage in a protocol that hides which credentials Bob's policy requires. The outcome of the blinded policy evaluation is that if Alice satisfies Bob's policy then she learns the requested message, and she learns nothing about the message if she does not satisfy Bob's policy. Due to the space limitations we cannot include the technical details of our protocols, but they will be in the full version of the paper. We now describe each phase in more detail:

- *Credential Hiding Phase:* There is some set of attributes $a_1, \ldots, a_n$ that are known to Bob (and possibly to Alice). At the end of this phase Alice has a set of values $k_1, \ldots, k_n$ (i.e., one for each credential), where $k_i \in \{r_i[0], r_i[1]\}$, which are random values generated by Bob. These values will either be encryption keys or seeds for a pseudo-random generator that can produce such keys. The value of k_i is subject to the following constraints: (1) $k_i = r_i[1]$ only if Alice has a credential satisfying a_i, otherwise she gets $k_i = r_i[0]$; and (2) A computationally-bounded Alice learns nothing about the value $\{r_i[0], r_i[1]\} - \{k_i\}$. Our full paper gives three protocols for this phase:
  1. *Protocol 1*: In this protocol it is assumed that Bob is willing to reveal to Alice a superset of the credentials in his policy. While this is not acceptable for all applications, there are many cases where Alice could guess with high probability the set of attributes in Bob's policy before the protocol, and in such cases this protocol may be acceptable to Bob. The communication complexity of this protocol is $O(\rho n)$ and it requires 3 rounds of interaction.
  2. *Protocol 2*: Unlike Protocol 1, this protocol does not assume that Bob is willing to reveal a superset of the attributes in his policy. In this protocol, Bob learns the value m and Alice learns: (1) the value n and (2) the number of attributes in Bob's policy that she satisfies (but she does not know which of her credentials satisfy Bob's attributes). This protocol requires $O(\rho mn)$ communication and 5 rounds of interaction.
  3. *Protocol 3*: This protocol is similar to Protocol 2, but Alice does not even learn how many attributes she satisfies in Bob's policy. This protocol requires $O(\rho^2 mn)$ communication and 3 rounds of interaction.
- *Blinded Policy Evaluation Phase:* Given the k values from the previous phase, Alice and Bob engage in a protocol that allows Alice to learn message M if she satisfies Bob's policy. His policy is represented by a boolean function $P : \{0,1\}^n \rightarrow \{0,1\}$ (i.e., it maps n values, which correspond to which attributes Alice has, to a binary value that corresponds to whether or not Alice satisfies Bob's policy or not). That is, if after the previous phase Alice's values are $r_1[x_1], r_2[x_2], \ldots, r_n[x_n]$, where $x_i \in \{0,1\}$, then Alice will receive M iff $P(x_1, x_2, \ldots, x_n) = 1$. Our work includes two protocols for blinded policy evaluation. The first protocol requires $O(2^n)$ communication for a policy with n attributes, and the second requires communication polynomial in n. However, the first protocol is for arbitrary functions and the second protocol is for a special class of functions.

4. CONCLUSION

We gave an efficient protocol for Alice to access a resource from Bob, such that Alice does not learn Bob's policy and Bob does not learn Alice's credentials. The only information Alice learns is whether she get access or not. Future work includes applying our protocol to trust negotiation systems.

5. ACKNOWLEDGEMENTS

We would like to thank the anonymous referees for their helpful comments.

6. REFERENCES

[1] D. Boneh and M. Franklin. Identity-Based Encryption from the Weil Pairing. In *Proceedings of Crypto 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 213–229. Springer, 2001.

[2] J. E. Holt, R. W. Bradshaw, K. E. Seamons, and H. Orman. Hidden credentials. In *Proceedings of the 2nd ACM Workshop on Privacy in the Electronic Society*, Oct. 2003.

[3] A. C. Yao. How to generate and exchange secrets. In *Proceedings of the 27th IEEE Symposium on Foundations of Computer Science*, pages 162–167. IEEE Computer Society Press, 1986.