Location-Based Access Control Systems for Mobile Users – Concepts and Research Directions

[Invited Paper]

Elisa Bertino
Purdue University
Department of Computer Science
West Lafayette, IN, USA
bertino@cs.purdue.edu

Michael S. Kirkpatrick
James Madison University
Department of Computer Science
Harrisonburg, VA, USA
kirkpams@jmu.edu

ABSTRACT

Many organizations require that sensitive information only be accessed on the organization premises or in secure locations. Access to certain information is thus allowed to authorized users, provided that these users are in specific locations when accessing the information. The GEO-RBAC model addresses such requirement. It is based on the notion of a spatial role, that is, a geographically bounded organizational function. The boundary of a role is defined as a geographical feature, such as a hospital or a classified facility; it specifies the spatial extent in which the user must be located in order to use the role. Besides a physical position obtained from a mobile terminal, users are assigned a logical and device independent position, representing the feature where the user is located. Logical positions are computed from real positions by specific mapping functions. If the user is present within the spatial boundary of a role, the role is said to be enabled. The user is allowed to select (activate) a role and exercise the associated permissions only once the role is enabled. The deployment of an access control system based on GEO-RBAC entails addressing several challenges: (1) access policies may require that access be conditioned not only by the user location but also on the presence or absence of other users; (2) enforcing location-based access control requires making the access control server aware of user locations, which may lead to privacy breaches; (3) trustworthy information about user locations must be obtained. This paper elaborates on these challenges and outlines related research directions.

Categories and Subject Descriptors

H.2.8 [**Database Management**]: Database Applications—*Spatial Databases and GIS*; K.4.1 [**Computers and Society**]: Public Policy Issues—*Privacy*

SPRINGL '11 November 1, 2011. Chicago, IL, USA

General Terms

Design, Languages, Security, Management, Verification

Keywords

Location-based access control, privacy

1. INTRODUCTION

Role-based access control (RBAC) [17] is a key industry standard for authorization that has been widely deployed across a variety of types of organizations. One of the main advantages of RBAC is a simplified mechanism for granting access to sensitive resources based on job functions, rather than identities. As mobile devices become more prevalent in enterprise settings, organizations and researchers have begun to explore techniques to enrich the model with spatial features [6, 8, 1, 2, 16, 4]. The Geo-spatial RBAC (GEO-RBAC) model was one of the first extensions to incorporate location constraints into access control policies, providing organizations with the ability to control resource access based on the physical coordinates of the requesting users.

The security requirements of military environments provide a natural scenario for deploying GEO-RBAC. For example, a traditional RBAC policy could protect surveillance data by restricting access only to the personnel acting in command roles. With GEO-RBAC, one could extend the policies to restricting this access to those personnel, while they are present in a secured command facility. Alternatively, a policy could require presence in a location that satisfies certain properties, such as certification that the room has been checked to ensure there are no unauthorized surveillance devices present. GEO-RBAC addresses these access control needs and supports permission manipulation at a fine granularity.

Military settings are not the only beneficiaries of the GEO-RBAC model. Consider the security requirements of a corporation developing a breakthrough technology. In order to prevent the company's research secrets from falling into the hands of a competitor, GEO-RBAC could be deployed to permit access to that data only from a corporate setting. Consequently, if an employee's laptop is stolen or lost, the access control system could prevent an adversary from recovering and selling the secret designs. Furthermore, some employees may not be trustworthy and pose a threat to the secrecy of the project. Enforcing spatial constraints could mitigate the threat that these insiders pose.

While location-based access control models, such as GEO-RBAC, are designed to address such stringent security requirements, it is desirable to extend the models with the following features: (1) a rich set of constraints that condition the access to sensitive resources to the presence or absence of other users in the same location of the user requesting access; (2) protection of user location privacy. In addition, integrating the access control system with technologies capable of providing trustworthy location data is crucial. In this paper we discuss preliminary approaches to address these requirements and outline research directions.

The rest of the paper is organized as follows. We first survey a proximity constraint language defined as part of our previous work [10]. We then discuss the problem of how to enforce location-based access control while at the same time assuring privacy. Finally we discuss an approach to obtain trusted location information, and then present some concluding remarks.

2. LOCATION PROXIMITY CONSTRAINT LANGUAGE

While researchers have developed a number of RBAC extensions that consider the requesting user's location, such protection models are not sufficient for all scenarios. That is, it is often more critical to consider the relative proximity of other users, specifically their presence or absence in the same space as the requester. For instance, consider a government official accessing a document that is classified according to a multi-level security model. Depending on the nature of the document, it may be desirable to require the presence of a supervisor, the absence of any civilian (*i.e.,* non-governmental personnel), or both.

As an initial approach into this new realm of spatially aware RBAC, we propose the definition of new extensions for specifying proximity constraints and on techniques for enforcing them. Such an approach entails addressing several issues, including:

- The identification of a suitable space model able to represent protected areas (PA). We define a PA to be a physically bounded region of space, accessible through a limited number of entry points, which consist of physical barriers that require authorization. Another critical issue is related to the definition and enforcement of a constraint language suitable for open spaces.
- The integration of a proximity language with GEO-RBAC. Such integration requires understanding at which level in GEO-RBAC the constraint language should be integrated: constraints could be specified at the level of role schemas, at the level of the role instances, or at both levels. The third approach appears the most flexible; however it requires addressing issues related to constraint overriding and refinement.
- Architectural approaches for access control enforcement. Such approaches depend on the user positioning and tracking technologies that are currently available; their integration with the architecture of the access control enforcement systems should be investigated.

We have proposed a language, called Prox-RBAC [10], for expressing such policies. Our approach starts with an indoor space model, in which the reference space is divided into hierarchical protected areas. For instance, multiple rooms exist as subspaces of suites, but the rooms are all non-overlapping. Using this space model, one can define *proximity constraints* based on the logical locations defined. That is, policies do not require physical coordinates; rather, they can be defined according to logical geospatial features, such as $Room205$ or $third\ floor$.

A proximity constraint is defined according to three primitive constructs: relative constraint clauses, continuity of usage constraints, and timeouts. First, *relative constraint clauses* define the static conditions that must be met in order for access to be granted. These constraints can be described as either absence or presence constraints. We adopt an intuitive syntax, illustrated as follows:

$$at\_most\ 0\ civilian\ in\ Room205$$

The core of the condition is the role (*e.g.,* civilian) and the spatial relationship (*e.g., in* $Room205$). The spatial relationship consists of a topological relation (*e.g., in*, *adjacent*) and a logical location descriptor for the space under consideration. While the location descriptor can be an explicit location as illustrated above, we also use the keyword *this* to define a relative location. For instance, *this.room* would require the constraint be evaluated relative to the room in which the requesting user is present. Note that this implies a typing relation on PAs for policy evaluation. In addition, optional cardinality qualifiers (*e.g., at_most n* or *at_least n*) can be added to facilitate flexible constraint definitions. Finally, as the number of conditions desired for a policy may be complex, Prox-RBAC allows relative constraints to be combined with boolean connectives.[1]

While this intuitive syntax captures the essence of a proximity constraint, it fails to address the issues that arise with mobility. That is, some operations may require a significant amount of time for completion, and the constraint must be satisfied for the entire duration. To express these requirements, we introduce the *when* and *while* keywords[2] to declare whether the condition must hold only prior to access or for the duration. Their use is illustrated as follows:

$$while\ (at\_most\ 0\ civilian\ in\ Room305)$$

In this example, the *while* constraint mandates that the policy enforcement point (PEP) continually evaluate the condition.[3] Given that the technology underlying the PEP is application dependent, the frequency of this check would also be a system parameter, requiring consideration of network latency, size of the spatial environment, number of users, and other such issues. On the other hand, a *when* constraint is only evaluated once when the access is requested. If the constraint is later violated (*e.g.,* a civilian entered the room), the access would not be revoked. As in the case of relative constraint clauses, continuity of usage clauses can also be combined using boolean connectives.

In many scenarios, it would be desirable for the policies to allow some flexibility in the granularity of the condition

[1]For practical deployment, we use boolean words, such as "*and*" and "*or*" instead of the symbolic notations $\wedge$ and $\vee$.

[2]As above, we adopt natural language keywords for our language. One could substitute a symbolic representation, such as linear temporal logic, as well.

[3]Observe that this implies the PEP takes on some of the duties of the policy decision point (PDP), as the PEP must decide if the condition still holds.

check. That is, there are times when the policy could be temporarily violated without revocation of a user's privileges. For instance, consider a policy that requires the presence of a supervisor while a document is being read. If the supervisor leaves the room for a short break and the policy is strictly enforced, the user's productivity may suffer, as his concentration is broken when the supervisor leaves; in such a case, it may be acceptable for the user to maintain the privileges while the supervisor is gone. To express this condition, every proximity constraint that includes at least one *while* clause may end with a timeout constraint, as follows. Note that, if the condition is satisfied prior to the timeout specification, the system behaves as if the constraint was never violated.

$$while\ (\ clause\ )\ timeout\ t$$

Future work in the area of proximity-based access control can focus on a number of important issues. First, techniques for lightweight continuous monitoring could be explored. In our prior work, we employed an *event-based* approach, in which any movement from one PA to another triggered a re-evaluation of the entire system. More efficient techniques that localize policy evaluation would be very desirable. Similarly, our work requires explicit actions by users to track the one-way nature of a move (*i.e.,* moving from PA_1 to PA_2 causes a different reaction than moving from PA_2 to PA_1). Streamlining the user experience for such a scheme would be very beneficial.

3. PRIVACY IN LOCATION-BASED ACCESS CONTROL

Existing models and enforcement architectures regarding location-based access control are generally built on a one-way security assumption. Specifically, these systems are designed to protect the sensitive data provided by the server against possible threats from remote users. However, these techniques offer no protection in the other way; the server-side is implicitly assumed to be trusted, and no user privacy controls are offered.

High-frequency disclosure of fine-grained location information can pose a severe privacy threat [7], especially when the server-side consists of a decentralized, loosely coupled environment such as cloud computing. In such settings, one cannot assume the server is fully trusted. Even when a centralized authorization infrastructure is in place, malicious insiders can still pose a threat to users. Consequently, it is crucial that these systems deploy measures that integrate privacy safeguards with security.

We have proposed a privacy-preserving enforcement architecture for GEO-RBAC [11] that allows the PEP to enforce its policies correctly, while preventing disclosure of the user's identity, role, or location. Our approach combines cryptographic techniques with a separation of functionality among several distinct components. The only trusted component in our architecture sets up the initial cryptographic data components, but does not participate in the on-line operations of the PEP. Thus, it can be effectively shielded from protocol-based attacks.

The back-end of our framework consists of four pieces: the service provider (*SP*), a role authority (*RA*), a location authority (*LA*), and an identity authority (*IA*). The *IA* is the trusted component that establishes the cryptographic secrets for the system. When a client wishes to make a request, it contacts the relevant *SP* that controls the resource, providing a pair of tokens. *SP* gets *RA* and *LA* to decrypt the tokens blindly through the use of commutative encryption. This process allows *SP* to authenticate that the tokens correspond to a valid role and location. *SP* then initiates an oblivious transer (OT) [15] session with *RA*, as well as private information retrieval (PIR) [5, 12] with *LA*, to retrieve additional data used for the policy evaluation. Prior to any request, each of *SP*'s policies are encoded as follows, where ρ indicates the role, λ the location, α the action, and δ the object, with system-wide constants ι and N:

$$< (\rho\lambda)^{-1}(\alpha\delta) \mod \phi(N), \iota^{\alpha\delta} \mod N >$$

SP has no knowledge of how to map this encoded policy to the original policy. When a client makes a request and provides tokens validating a claim to a role and location, *SP* can determine if one of the policies are satisfied, but never learns what the original, unencoded policy was. Thus, *SP* cannot learn the user's role or location. For further details, including proofs of the security properties, please see [11].

We have proposed three directions for future research in this area. First, our approach is similar to attribute-based cryptography, but our credentials are inherently transient; we propose exploring the application of this concept to other domains. Next, we raise the question of other settings that can reconcile security and privacy. Third, our scheme is vulnerable to attack if *SP* is able to collude with either *RA* or *LA*; techniques that can mitigate this threat would be of great interest.

4. TECHNIQUES FOR TRUSTWORTHY LOCATION INFORMATION

An access control system with location-based constraints requires some mechanisms for securely determining the user location and tracking user movements. Today several technologies exist for such purposes, including sensor-based user tracking. One can also combine location-based RBAC with physical access control, through keypad locks on doors or smartcard. Additionally, one can integrate these techniques with systems for digital identity management to authenticate a user's location claim with high levels of assurance. The definition of techniques for secure location identification entails addressing several issues, including:

- **Usability.** As described in [3], it is important that any actions a user must perform should be minimized and intuitive.
- **Strong location assurance.** Depending on the type of device and technologies used for sensing the user location, physical threats, like the device being stolen or voluntarily released by the owner to other users, are possible. In such case we need techniques to ensure that the device is actually used by his owner.
- **Availability.** In emergency situations, access control policies may have to be by-passed to allow the user to access resources needed to manage the emergency. Health care is an application domain example that has such a requirement. We thus need approaches by which access control policies may be allowed to be violated; however if a violation occurs, the system must

collect additional evidence that can be user later to audit the violation and determine whether the bypass was justified. In the case of mobile devices it is important to augment the proposed access control system with mechanisms for collecting evidence related to location and movements, for example evidence that a user has been the first one to arrive at a location and nobody else was present.

An approach that we have proposed is based on the use of Near Field Communication (NFC) technologies [9]. NFC is an RFID-based proximity-constrained technology that provides contactless communication between a device and a reader/writer. In contrast to traditional one-way RFID mechanisms, NFC has a number of advantages. NFC has a very restricted broadcast range, typically 10 cm, which helps to create a stronger assurance of the user's location. NFC also defines a peer-to-peer mode that exchanges data in both directions in a single contactless session. This mode mitigates the threat of passive attacks that steal data stored in NDEF tags [13, 14]. As our approach does not use NDEF tags, we avoid these threats.

Unfortunately, the trust assumptions in our scheme are very strong, requiring secure storage of identity information within the cell phone. Such an approach, in which a user is bound to a single trusted device, has negative implications on all of these desirable goals. Addressing these shortcomings is an important step toward practical application of these access control models.

5. CONCLUDING REMARKS

In this work, we have described a number of the challenges and concepts that exist within the realm of location-based access control. We have summarized our existing work in the field, including a policy language for considering other user's locations, cryptographic techniques for protecting users' privacy, and an enforcement architecture and implementation for secure location authentication. In addition, we have proposed a number of directions for future research in the field. While the incorporation of location information into access control policy models has been extensively explored, we argue that addressing these other technical challenges will be important for adoption of these models in practical settings.

6. ACKNOWLEDGEMENTS

The work reported in this paper has been partially supported by Sypris Electronics and by the MURI award FA9550-08-1-0265 from the Air Force Office of Scientific Research.

7. REFERENCES

[1] S. Aich, S. Sural, and A. K. Majumdar. STARBAC: Spatiotemporal role based access control. In *OTM Conferences*, 2007.

[2] C. A. Ardagna, M. Cremonini, E. Damiani, S. De Capitani di Vimercati, and P. Samarati. Supporting location-based conditions in access control policies. In *Proc. of 1st ACM Symposium on Information, Computer and Communications Security (ASIACCS)*, 2006.

[3] L. Bauer, L. F. Cranor, M. K. Reiter, and K. Vaniea. Lessons learned from the deployment of a smartphone-based access-control system. In *Proc. of 3rd Symposium on Usable Privacy and Security (SOUPS)*, 2007.

[4] S. M. Chandran and J. B. D. Joshi. LoT RBAC: A location and time-based RBAC model. In *Proc. of 6th International Conference on Web Information Systems Engineering (WISE '05)*, pages 361–375, 2005.

[5] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan. Private information retrieval. *J. ACM*, 45:965–981, November 1998.

[6] M. L. Damiani, E. Bertino, B. Catania, and P. Perlasca. GEO-RBAC: A spatially aware RBAC. *ACM Transactions on Information and System Security (TISSEC)*, 2007.

[7] G. Ghinita, P. Kalnis, A. Khoshgozaran, C. Shahabi, and K.-L. Tan. Private queries in location based services: anonymizers are not necessary. In *Proc. of 2008 ACM SIGMOD International Conference on Management of Data*, pages 121–132, 2008.

[8] F. Hansen and V. Oleschuk. SRBAC: A spatial role-based access control model for mobile systems. In *Proc. of 8th Nordic Workshop on Secure IT Systems (NORDSEC '03)*, pages 129–141, October 2003.

[9] M. S. Kirkpatrick and E. Bertino. Enforcing spatial constraints for mobile RBAC systems. In *Proc. of 15th ACM Symposium on Access Control Models and Technologies (SACMAT '10)*, 2010.

[10] M. S. Kirkpatrick, M. L. Damiani, and E. Bertino. Prox-RBAC: A spatially aware proximity-based RBAC. In *Proc. of 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM SIGSPATIAL GIS)*, 2011.

[11] M. S. Kirkpatrick, G. Ghinita, and E. Bertino. Privacy-preserving enforcement of spatially aware RBAC. *Transactions on Dependable and Secure Computing (TDSC)*, January 2012.

[12] E. Kushilevitz and R. Ostrovsky. Replication is not needed: single database, computationally-private information retrieval. In *Proc. of 38th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 364–373, 1997.

[13] G. Madlmayr, J. Langer, C. Kantner, and J. Scharinger. NFC devices: Security and privacy. In *The 3rd International Conference on Availability, Reliability and Security (ARES)*, pages 642–647, 2008.

[14] C. Mulliner. Attacking NFC mobile phones. In *EUSecWest*, May 2008.

[15] M. O. Rabin. How to exchange secrets with oblivious transfer. Technical Report TR-81, Aiken Computation Lab, Harvard University, 1981.

[16] I. Ray, M. Kumar, and L. Yu. LRBAC: A location-aware role-based access control model. In *Proc. of International Conference on Information Systems Security (ICISS)*, pages 147–161, 2006.

[17] R. Sandhu. Role-based access control models. *IEEE Computer*, February 1996.