Context-Dependent Authentication and Access Control

Michael Kirkpatrick[1] and Elisa Bertino[2]

[1] Department of Computer Science and CERIAS,
Purdue University, West Lafayette, IN, USA,
mkirkpat@cs.purdue.edu
[2] Department of Computer Science and CERIAS,
Purdue University, West Lafayette, IN, USA,
bertino@cerias.purdue.edu

Abstract. As mobile computing continues to rise, users are increasingly able to connect to remote services from a wide range of settings. To provide this flexibility, security policies must be adaptive to the user's environment when the request is made. In our work, we define context to include the spatiotemporal aspects of the user request, in addition to quantifiable environmental factors determined by the server hosting the resource. We identify a number of key open problems in this field and propose potential solutions to some of the problems.

1 Introduction

As computer networks, both fixed and wireless, expand their capacities and coverage, users are increasingly able to connect to remote services from a wide range of settings. An employee could be using a VPN connection from a home office. A student could initiate an SSH session using a public wireless network at a cafe. Someone could check his bank account balance using a cell phone or PDA. All of these examples show that the possibility of connecting from any location is crucial for user convenience. However, it also introduces the need for more articulated security requirements. In this paper, we focus on one such requirement, which is the problem of incorporating contextual information into access control decisions.

Traditional approaches to access control have been built on a static view of computing. That is, if the subject is granted access to a resource, he can typically exercise this permission from any setting. For example, discretionary access control (DAC) systems grant subjects the ability to perform actions based solely on an authenticated identity. If the user can login to the system using his username and password, he can access his files, regardless of whether the connection is performed via a secure shell or an unencrypted telnet session. As DAC is based only on the identity authentication, it clearly does not provide the ability to restrict access based on which machine is used, what time the access request is sent, or what type of network connection is used.

Mandatory access control (MAC) addresses a lot of the shortcomings of DAC by introducing system-enforced constraints. Under DAC, the owner of an object or file makes the decisions of granting permissions to other users. MAC systems place additional constraints that are beyond the control of the object owner. For example, in multi-level security (MLS), subjects and objects are both tagged with security classifications; a subject's request to access an object is granted only if the classification hierarchy allows it. Although MAC policies place restrictions beyond the identity authentication of DAC, MAC also does not consider the context of the request.

Role-based access control (RBAC) is another widely used approach. RBAC is similar to DAC, except the notion of identity is replaced with a role. For instance, in a health care setting, permissions for a patient's records can be restricted to the role *Doctor*; any user who is then capable of activating the *Doctor* role would then be granted access to the record. RBAC is similar to the notion of *group* in UNIX file systems. RBAC is more powerful, though, as it can be used to enforce policy constraints such as separation of duty. However, as with MAC and DAC, RBAC does not provide adequate functionality to incorporate and adapt to contextual information.

In this work, we explore novel approaches to access control that attempt to create a more adaptive, fine-grained systems by considering contextual information. We describe this area of work as context-dependent authentication and access control (CDAC). Our work is primarily motivated by the needs of organizations to provide flexible mobile access to registered users. Examples of such organizations include enterprises with remote employees, law enforcement agencies with mobile personnel responding to emergencies, and health care networks with geographically distributed sites, as well as mobile staff and physicians. Our work is also relevant to the problem of secure, ad hoc information sharing, sometimes referred to as the dynamic coalition problem. In such a setting, businesses or governments cooperate by sharing information to accomplish a short-term goal, despite the possibility of conflicting long-term aims.

An overarching theme of our work is the lack of focus on the problems of CDAC enforcement. Study of access control has traditionally split the topic into the dichotomy of policy and implementation, with the implicit view that policy is a subject for security research, while implementation issues are simply questions of engineering and less deserving of research consideration. Sandhu *et al.* [38] argued that the separation between policy and implementation in advanced access control systems is too great, and more work needs to focus on how to bridge this gap. They proposed the concept of PEI (policy, enforcement, and implementation) models to emphasize the need for focus on issues of enforcement. We agree with this approach and, as such, examine the need for future research in the area of enforcement.

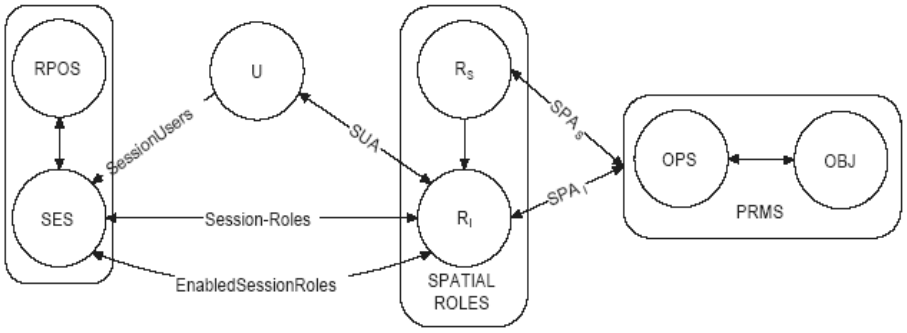

Fig. 1. Core features of GEO-RBAC

2 Background

The first challenge in CDAC is to define precisely what is considered context. Intuition dictates that context should reflect the user's environmental conditions. Clearly, the user's physical location, represented by GPS coordinates, could be considered an example of context. In some settings, the logical location may be more useful; that is, the user's location is described in relative terms, such as "on the third floor," "in the hospital emergency room," or "in room 217." The precise GPS coordinates may be unnecessary for access control in such settings. Several models [13, 24, 1, 36, 7] have been proposed to incorporate the user's geographic location into role-based access control. Additional work in CDAC based on location includes [23], which focuses on privacy-preservation in location-based services.

GEO-RBAC [13] introduces the concept of *spatial roles*, combining a traditional RBAC role with particular spatial extents. During a single session, a user is mapped to one or more spatial roles according to his or her location, as well as any credentials required to activate a role. *Permissions*, linking operations and objects that can be acted upon, are assigned to spatial roles. Thus, if a user can activate a particular spatial role during a session, he or she can then perform the actions specified by that role's permissions. The core notions of GEO-RBAC are shown in Figure 1.

There are two key novel features to GEO-RBAC. The first is the distinction between *role enabling* and *role activation*. When a user enters the region described by the spatial role's extents, we say that the role is *enabled*. However, the user cannot exercise any permissions associated with that role until he chooses to *activate* it. If the role is not activated, the user cannot exercise any of the associated permissions. The advantage of this distinction is that mutually

exclusive roles can be defined for the same spatial region. Both roles can be simultaneously enabled, but only one can be activated at a time.

Another key feature of GEO-RBAC is the concept of *role schema*. A role schema is an abstraction, such as $<$ *Doctor*, *Hospital* $>$, that can be used as a template for singular roles. Permissions can be granted to role schemas in order to ease the administration of roles. A *role instance* is then created from a schema using specific data. For example, $<$ *Chief of Surgery*, *St. Vincent* $>$ can be an instance created from the $<$ *Doctor*, *Hospital* $>$ schema, as *Chief of Surgery* is a particular instance of *Doctor*, and *St. Vincent* is an instance of *Hospital*. Note that, while permissions can be granted to either schemas or instances, only instances can be activated. For additional details on GEO-RBAC, we refer the reader to the full paper.

Location is not the only factor commonly used as contextual information. Some models incorporate the time of the access request [1, 5, 3]. For example, an organization may wish to restrict access to sensitive data to business hours. Additionally, more subtle factors can be considered. The user's previous data access history may be required to enforce separation of duty or conflict of interest constraints. Environmental factors, such as the presence of a medical emergency or a criminal pursuit, may be sufficient to grant an access request that would otherwise be denied. Previous work in CDAC [30, 26] also includes aspects such as velocity or physical world conditions. Finally, context-awareness has also been applied to the unique challenges of ubiquitous computing in the home [10].

Another aspect of contextual information for consideration is the trustworthiness of the principal. In order to quantify this factor, researchers have focused on the calculation of either risk or trust [8, 14, 15, 4]. One challenge in this field is to design the system to adapt to new information. For example, many risk-based approaches involve defining weighting factors for pieces of data. However, it is not clear how a system should react when a user presents a new form of credential. Furthermore, as we will describe in Section 7, risk- and trust-based approaches may actually work against the goals of the access control system.

Our definition of *context* reflects the sum of when, where and how the user makes his request, in addition to certain environmental conditions that are beyond the knowledge or control of the user. That is, we define context to be the combination of quantifiable data that may be relevant to an access control decision. This definition includes (but is not limited to) the user's spatiotemporal setting, his access request history, the device used to make the request, the trust placed in the user by the organization, the time of the access, the frequency of access requests, and the presence of an emergency situation (as reported by the user and/or other trustworthy sources). We emphasize that a suitable notion of context includes many different aspects, thus making the design and implementation of a CDAC system challenging.

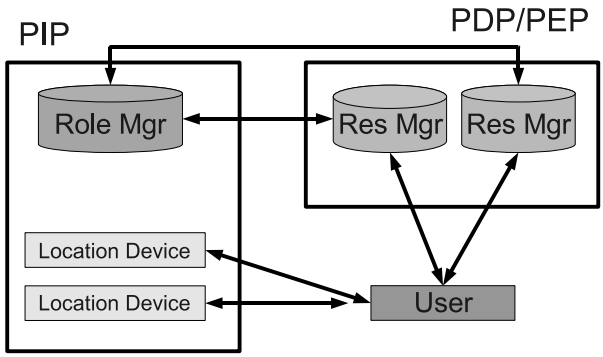

Fig. 2. Architecture for spatially-aware RBAC with continuity of usage

3 Context Determination

Once the notion of context has been defined, system designers must address the question of how to determine the true context of a request. In some settings, it may be appropriate to trust the user's device to report the contextual information reliably. For example, in spatiotemporal access control, certified software on a portable device may be able to use GPS coordinates directly in determining whther to grant an authorization.

We have a current project that explores the use of CellDB [6] for such an implementation. Most cell phones define an API that allows a program to determine the unique cell ID based on the user's location. The program can then submit a query to CellDB to convert the cell ID to a set of coordinates that estimate the GPS coordinates with a high degree of accuracy. CellDB is built on the same technology that underlies Google Latitude [19].

While capabilities such as CellDB and GPS are sometimes helpful, in a distributed setting, the policy enforcement point (PEP) may exist separate from the user's device. In such an environment, the system must ensure that the reported location is indeed correct. That is, the server should not necessarily accept the device's claim to a location. The system must contain an enforcement mechanism that authenticates the claim.

One method for enforcing location constraints is to use a number of devices, such as a pre-deployed wireless sensor network, placed in physically distinct locations. In such a design, the sensors can detect the user's device and report the location reliably to other nodes, relaying the location to master nodes or servers that enforce the policies. In [29], we have explored a different technique based on contactless proximity devices, such as cell phones enabled with Near-Field Communication (NFC) technology [31, 32].

Figure 2 shows a high-level view of our architecture. To initiate an access request, the user device obtains a proof of location from the pre-deployed location device. The proof includes a timestamp to prevent replay. Furthermore, NFC

technology has a limited broadcast range, so the proof guarantees the user's proximity to the device. The proof is then sent to the resource manager (Res Mgr), which serves as both the policy decision point (PDP) and policy enforcement point (PEP). In order to make a proper authorization decision, the resource manager forwards the user's credentials, along with the proof of location, to the central role manager (Role Mgr). The role manager then returns a list of authorized roles, which the resource manager considers in order to grant or deny access. Note that, as the role manager and location device both provide relevant information, we consider both to be a part of the policy information point (PIP).

In addition to location, other contextual factors pose their own authentication challenges. For example, if a police officer is attempting to gain access to a building's floorplan in pursuit of a suspected criminal, it would be desirable for the system to trust the officer's judgment. However, if the officer's device has been compromised, the claim of an emergency may be false as part of an attempt to gain illicit access. Consequently, authentication of quantifiable environmental factors should also be considered as part of the system design process. To our knowledge, little focus has been placed on this issue. One reason for the lack of attention is the diversity of environmental factors that can be considered. However, exploration of this subject warrants further consideration.

Physical authentication of the user device poses additional problems for CDAC. Attestation schemes [37, 39] and approaches to secure roaming [25] can be used to ensure that a remote (untrusted) device guarantees certain behavior using trusted hardware. Trusted hardware can also be used to store cryptographic keys or bind protected data to a machine [42, 2, 18, 22, 16, 21]. It would be interesting to consider the integration of these hardware approaches into the authentication process. That is, instead of using the trusted hardware to secure a key used for encryption, the module produces a unique value that is used in addition to other authentication factors.

We have proposed one such approach [28]. Our intended application is for protecting data of multiple levels of sensitivity across physically remote branches. At each site, there are a number of workstations that are administered to be more secure than other machines. The most sensitive data must be accessed only by these workstations. Furthermore, as the application may need access to data of multiple levels of protection, the access control decision must be made at the application layer, rather than the network layer. I.e., solutions based on VPNs, IPsec, or DNSsec are not appropriate.

In addition, our approach was designed with the goal of defending against the insider threat. That is, no single administrator could authorize a new machine to access the sensitive data. We used a k-of-n secret splitting approach so that multiple administrators were required for the setup process. Once the key was installed in the trusted hardware, it could not be leaked and physical access to the machine was required (though not sufficient) for authorization to use the sensitive data. We refer the reader to [28] for further details.

Other contextual factors may introduce their own authentication problems. While there are many works that propose the usage of contextual information

into access control, these works do not address the problem of how to ensure the correctness of this data. We emphasize here the need for additional work to solve these problems and to demonstrate the feasibility of designing systems that can ensure the veracity of such contextual information.

4 Usage of Context

Assuming context has been defined and the system has performed any necessary check to ensure the validity of the claims of the context, there is a fundamental issue that is frequently overlooked. Specifically, the system designer must make a decision as to how knowledge of the context is used. The vast majority of work in CDAC assumes the information is used in the style of multi-factor authentication. For example, location-based models, such as GEO-RBAC and STARBAC, expand the notion of *role* to encompass a spatial region; thus, if the user can show proof of location and can provide adequate credentials to activate the role, the access control policy is satisfied and authorization is granted.

While the previous usage is perhaps the most intuitive, it is not the sole approach. Another technique is to use the context information *to determine the policy*. In this view, the context is distinct from the access control decision process. Rather, it is a requisite pre-cursor to the decision. As an example, consider the needs of a consulting business with mobile employees. These employees work from different locations with different levels of security. When deployed for an assignment, they may be working at another company's office building. In between assignments (occasionally referred to as "on the bench"), the employee may come into his company's local office. At other times, the employee may be connecting from home or from an unsecured wireless network at a café.

A naive approach would be to use solely a VPN connection from remote locations. The advantage of a VPN is clear, as any data transmitted between the user's machine and the company's servers would be encrypted. However, more security would be desirable as a result of the threats that emerge from different networks. For example, if the current network is not adequately administered, the user's machine could become compromised by malicious software. The VPN would protect the transmission of the data to the machine, but the malware may then leak the data to an external source. Consequently, it would be advantageous if the employee's company's policy's restricted access to sensitive data according to the type of network connection. A conservative approach, in which such requests are always denied, may be undesirable, especially in time-critical circumstances. A better solution would be to use the context to determine the applicable policies.

Auth-SL [41] defines a policy language for crafting multiple authentication policies. For example, these policies can dictate that a user must present more or stronger credentials for sensitive data. The possibility of applying the same technique to CDAC has been unexplored. Consider the case of a dynamic coalition in which the cooperating agencies do not fully trust one another. One organization may wish to monitor the requests from the other. If the other agency is

issuing a continuous stream of requests for sensitive data (some of which may not be relevant to the stated cooperative aim), it would be good to apply a more restrictive set of policies; if the other agency is only requesting a small amount of data that is directly related to the joint mission, the same restrictions may not be necessary. In short, we see the possibility of a CDAC mechanism that performs a risk analysis to determine the requisite security policies according to the context.

5 Enforcement Models

As previously described, most of the work in CDAC has involved defining abstract models for reasoning about environments and contexts. Bridging the gap between these models and the development of real systems can be a daunting task. As such, we would like to see more attention on defining generic frameworks for the development of CDAC systems. XACML [33] is a good example of a framework that can help in the design of access control systems. It defines concepts such as the PEP, PDP, and PIP, among others.

For CDAC, it would be advantageous to have a similar reference standard for handling contextual data. Previous works have proposed designs under particular assumptions, for instance, requiring a pre-deployed sensor network or a homogeneous environment in which all users' devices achieve and report a consensus of the current settings. It is an open question as to whether one could define abstract principals and generic protocols that can be applied without such requirements.

One topic regarding enforcement models that requires consideration is the timing of access control checks. Traditional approaches to authorization is to evaluate the user's request prior to granting access, verifying the presented credentials and applying the policy only at that time. Many previous works in CDAC have adopted the same methodology. If the user can present the required proof of identity while certain contextual constraints hold, the request is granted.

The problem with this traditional method is that it assumes a static view of the environment, whereas CDAC systems are inherently dynamic. For example, in a spatiotemporal authorization mechanism, the system should assume that the user is mobile, and may leave the authorized region. That is, CDAC designs must address the problem of continuity of usage [46, 34, 12, 45, 44]. The dynamic nature of contextual data requires that access decisions must be repeatedly evaluated even after the initial request is granted.

The definition of an *event-based* approach to continuity of usage for CDAC systems would be particularly useful. That is, if the context changes sufficiently during usage of a resource, the system would trigger an *event* that forced re-evaluation of the user's authorizations. In Figure 3, we illustrate events that can be triggered by a mobile user in a location-aware system. Event e_1 occurs when the user enters region A, and e_2 is triggered when crossing into region B. Finally, e_3 indicates the user has left A, but is still within the extents of B.

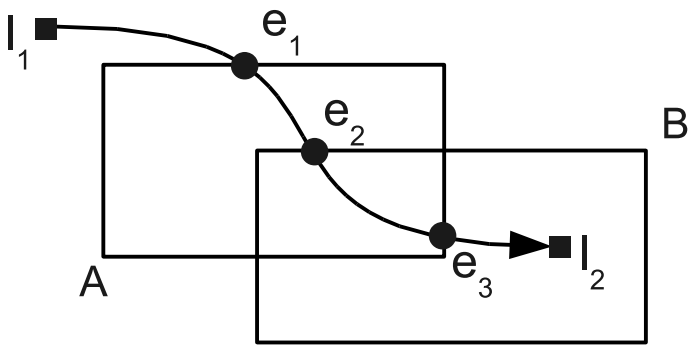

Fig. 3. As the user follows the path from location l_1 to l_2, events e_1, e_2, and e_3 are triggered when the user enters or exits regions A and B

Note that any of these events can either add or remove authorizations. For example, mutual exclusivity constraints may indicate that the user must be within A but not within B. In such a case, the user should lose some authorized capabilities after event e_2. Similarly, event e_3 can either yield more permissions (assuming the requirement that the user must be in B but not A) or it can result in dropped permissions. In any case, it is the *change in context* as the user crosses regional boundaries that triggers the event to force re-evaluation of the applicable permissions. Unfortunately, we are aware of no such technique for event-based continuity of usage in CDAC.

6 Partial or Conflicting Information

One topic for research in CDAC with no clear direction is how to make access decisions given partial or conflicting context information. Approaches such as [14] require *a priori* definition of weighting factors for contextual issues. One flaw with this approach is that it cannot automate the addition of new types of data. That is, if the user presents a new type of credential that does not have a weighting factor assigned, manual intervention is required. Additionally, there is no clear mechanism that handle contradictory evidence.

Alternatively, models based on fuzzy logic are very adaptable and require no *a priori* knowledge of the contextual factors. In FuzzyMLS [8], there is not a single barrier between granting and denying a request. Rather, there are a number of degrees between the two choices. The challenge with this approach, though, is that the targeted model requires manual intervention. That is, FuzzyMLS is not intended as an automated process, but rather it provides risk quantification that a domain expert can consult to make a decision.

It is an open question of whether a CDAC system can be designed to react to partial information. Our desired features for such a mechanism are that it

can be automated, it can adapt to new information without requiring *a priori* knowledge of the context factors, and it can handle contradicatory reports of the context. Traditional probabilistic approaches, such as Bayesian analysis, may be difficult to apply, as the probabilities of events may be unknown. One possible alternative would be to apply Dempster-Shafer theory, which constructs probabilities from evidence [40]. While Dempster-Shafer models appear to have applicable mathematical structures for conflicting information, we are unaware of any attempt to incorporate these models into a computing or access control system. As such, the performance of such a scheme is entirely unknown.

An additional possibility would be for the system to provide a partial response. For example, assume the user submits a number of credentials along with a set of requests, but the PDP can only validate a subset of the credentials (possibly due to network outages or revocations). It may be possible to grant the subset of the requests that correspond to the valid credentials. We would be interested to see if techniques from probability, decision theory, or machine learning could be applied to these challenges.

7 The Paradox of Trust

Finally, we close with a discussion that we call the *paradox of trust.* As described previously, CDAC aims to increase the capabilities of mobile users without sacrificing requisite protections of data. Traditional approaches to access control are simply incapable of adequately expressing the *intent* of a security policy that can adapt to novel situations. Incorporating trust and risk into the access control mechanism, therefore, seems necessary to make the system flexible enough to provide security guarantees without interfering with legitimate work processes. However, these approaches may lead to unintended consequences.

Fundamentally, access control is a mechanism for restricting the actions allowed for *authorized users.* Consequently, access control inherently incorporates a *distrust* of the user. If the user was fully trusted to behave properly, then access control would be unnecessary. The user would self-govern his requests.

Allowing more flexibility by designing a system to trust the user based on contextual information, then, goes directly against the original goal of access control. Specifically, trust- and risk-based systems create a larger vulnerability against insider threats, which remain a real and underestimated problem [9, 27, 11, 43, 20, 35]. As automated systems are designed to grant more flexibility to *seemingly* trustworthy insiders, the likelihood of success for an insider-based attack increases. Furthermore, as a user gains access to more sensitive data, the profitability of a violation increases, as does the temptation. The paradox of trust, then, reflects the notion that *incorporating trust* can result in making the security mechanism *less trustworthy.*

Recent work has proposed the use of Bayesian techniques to detect insider threats. The Intelligent Insider Threat Detection (I^2TD) [17] system is one such example. While intrusion detection systems focus on determining whether or not an outsider has penetrated the defenses of a network, I^2TD uses anomaly

detection to determine if an *insider* has acted maliciously. However, this work is still in the preliminary stages of development. We would like to see more work in this area, as we feel that protections against insider threats will continue to grow in importance.

8 Conclusions

In this work, we have defined *context-dependent authentication and access control (CDAC)* as a means of considering quantifiable environmental conditions in the field of access control. We have identified the key problems of context authentication, usage of context, continuity of usage, and the incorporation of risk and trust, among other research areas. In each of these topics, we have identified a number of challenges and directions for future research. We argue that solving these problems will lead to mechanisms that grant more flexibility for the users, while preserving greater security than systems such as DAC. These mechanisms can be used to improve security in overlay networks and distributed systems.

9 Acknowledgement

The material reported in this paper is based in part upon work supported by the National Science Foundation under grant 0430274, and by the MURI award FA955-08-1-0265 from the Air Force Office of Scientific Research.

References

1. Aich, S., Sural, S., Majumdar, A.K.: "STARBAC: Spatiotemporal Role Based Access Control." *OTM Conferences*, (2007).
2. Atallah, M.J., Bryant, E.D., Korb, J.T., Rice, J.R.: "Binding Software to Specific Native Hardware in a VM Environment: The PUF Challenge and Opportunity." *VMSEC '08*, (2008).
3. Atluri, V., Chun, S.A.: "A Geotemporal Role-Based Authorisation System." *International Journal of Information and Computer Security*, vol. 1, 143–168 (2007).
4. Aziz, B., Foley, S.N., Herbert, J., Swart, G.: "Reconfiguring Role Based Access Control Policies Using Risk Semantics." *Journal of High Speed Networks*, Special issue on Security Policy Management 15, no. 3, 261–273 (2006).
5. Bertino, E., Bettini, C., Samarati, P.: "A Temporal Authorization Model." *ACM Conference on Computer and Communications Security (CCS '94)* (1994).
6. CellDB `http://www.celldb.org/`.
7. Chandran, S., Joshi, J.: "LoT RBAC: A Location and Time-Based RBAC Model." *Proceedings of the 6th Internation Conference on Web Information Systems Engineering (WISE '05)*, 361–375 (2005).
8. Cheng, P.-C., Rohatgi, P., Keser, C.: "Fuzzy MLS: An Experiment on Quantified Risk-Adaptive Access Control." *DIMACS Workshop on Information Security Economics* (2007).
9. Chinchani, R., Iyer, A., Ngo, H.Q., Upadhyaya, S.: "Towards a Theory of Insider Threat Assessment." *International Conference on Dependable Systems and Networks (DSN '05)* (2005).

10. Covington, M.J., Long, W., Srinivasan, S., Dev, A.K., Ahamad, M., Abowd, G.D.: "Securing Context-Aware Applications Using Environment Roles." *Proceedings of the 6th ACM Symposium on Access Control Models and Technologies (SACMAT)*, 10–20 (2001).
11. CSO Magazine and CERT and United States Secret Service: "2004 E-Crime Watch Survey: Summary of Findings." `http://www.cert.org/archive/pdf/2004eCrimeWatchSummary.pdf` (2004).
12. Damiani, M.L., Bertino, E.: "Access Control and Privacy in Location-Aware Services for Mobile Organizations." *7th Internation Conference on Mobile Data Management* (2006).
13. Damiani, M.L., Bertino, E., Catania, B., Perlasca, P.: "GEO-RBAC: A Spatially Aware RBAC." *ACM Transactions on Information Systems and Security* 10, no. 1 (2007).
14. Diep, N.N., Hung, L.X., Zhung, Y., Lee, S., Lee, Y.-K., Lee, H.: "Enforcing Access Control Using Risk Assessment." *Proceedings of the Fourth European Conference on Universal Multiservice Networks (ECUMN)*, 419–424 (2007).
15. Dimmock, N., Belokosztolszki, A., Eyers, D., Bacon, J., Moody, K.: "Using Trust and Risk in Role-Based Access Control Policies." *Proceedings of the 9th ACM Symposium on Access Control Models and Technologies (SACMAT)* (2004).
16. Dyer, J.G., Lindemann, M., Perez, R., Sailer, R., van Doorn, L., Smith, S.W., Weingart, S.: "Building the IBM 4758 Secure Coprocessor." *IEEE Computer* 34, no. 10, 57–66 (2001).
17. Ferragut, E., Sheldon, F., Neergaard, M.: "ITD (Insider Threat Detection) System." Oak Ridge National Laboratory (ORNL) Cyberspace Sciences & Information Intelligence Research (CSIIR) Group `http://www.ioc.ornl.gov/documents/factsheets/ITD%20Insider%20Threat%20Detection%20FactSheet.pdf`.
18. Gassend, B., Clarke, D., van Dijk, M., Devadas, S.: "Controlled Physical Random Functions." *Proceedings of the 18th Annual Computer Security Applications Conference (ACSAC)* (2002).
19. Google Latitude `http://www.google.com/latitude/`
20. Greitzer, F.L., Moore, A.P., Cappelli, D.M., Andrews, D.H., Carroll, L.A., Hull, T.D.: "Combating the Insider Cyber Threat." *IEEE Security and Privacy* 6, no. 1, 61–64 (2008).
21. Guajardo, J., Kumar, S.S., Schrijen, G.-J., Tuyls, P.: "FPGA Intrinsic PUFs and Their Use for IP Protection." *Proceedings of the 9th Cryptographic Hardware and Embedded Systems Workshop (CHES)*, 63–80 (2007).
22. Guajardo, J., Kumar, S.S., Schrijen, G.-J., Tuyls, P.: "Physical Unclonable Functions and Public-Key Crypto for FPGA IP Protection." *International Conference on Field Programmable Logic and Applications*, 189–195 (2007).
23. Han, K., Kim, K.: "Enhancing Privacy and Authentication for Location Based Service using Trusted Authority." *2nd Joint Workshop on Information Security* (2007).
24. Hansen, F., Oleschuk, V.: "SRBAC: A Spatial Role-Based Access Control Model for Mobile Systems." *Proceedings of the 8th Nordic Workshop on Secure IT Systems (NORDSEC '03)*, 129–141 (2003).
25. Hoang, L.N., Laitinen, P., Asokan, N.: "Secure Roaming with Identity Metasystems." *IDtrust '08* (2008).
26. Hulsebosch, R.J., Salden A.H., Bargh M.S., Ebben P.W.G., Reitsma, J.: "Context Sensitive Access Control." *Proceedings of the 10th Symposium on Access Control Models and Technologies (SACMAT)*, 111–119 (2005).
27. INFOSEC Research Council (IRC): "Hard Problem List." Department of Homeland Security Cyber Security Research & Development Center (2005).

28. Kirkpatrick, M., Bertino, E.: "Physically Restricted Authentication with Trusted Hardware." Under submission (2009).
29. Kirkpatrick, M., Bertino, E.: "An Architecture for Spatially-Aware RBAC with Continuity of Usage." Under submission (2009).
30. Kulkarni, D., Tripathi, A.: "Context-Aware Role-based Access Control in Pervasive Computing Systems." *Proceedings of the 13th Symposium on Access Control Models and Technologies (SACMAT)* (2008).
31. NFC Forum Tag Type Technical Specifications `http://www.nfc-forum.org/`.
32. Nokia 6131 NFC SDK Programmer's Guide.
33. Organization for the Advancement of Structured Information Standards (OASIS): eXtensible Access Control Markup Language (XACML) `http://www.oasis-open.org/committees/tc\_home.php?wg_abbrev=xacml/`.
34. Park, J., Sandhu, R.: "The UCON$_{ABC}$ Usage Control Model." *ACM Transactions on Information and System Security* 7, no. 1, 128–174 (2004).
35. Predd, J., Pfleeger, S.L., Hunker, J., Bulford, C.: "Insiders Behaving Badly." *IEEE Security and Privacy* 6, no. 4, 66–70 (2008).
36. Ray, I., Kumar, M., Yu, L.: "LRBAC: A Location-Aware Role-Based Access Control Model." *Proceedings of Internation Conference on Information Systems Security (ICISS)*, 147–161 (2006).
37. Sailer, R., Jaeger, T., Zhang, X., van Doorn, L.: "Attestation-based Policy Enforcement for Remote Access." In *Proceedings of the 11th ACM Conference on Computer and Communications Security (CCS '04)*, 308–317 (2004).
38. Sandhu, R., Ranganathan, K., Zhang, X.: "Secure Information Sharing Enabled by Trusted Computing and PEI Models." *ASIACCS '06: Proceedings of the 2006 ACM Symposium on Information, Computer and Communications Security*, 2–12 (2006).
39. Schellekens, D., Wyseur, B., Preneel, B.: "Remote Attestation on Legacy Operating Systems With Trusted Platform Modules." *Science of Computer Programming*, 13–22 (2008).
40. Sentz, K., Ferson, S.: "Combination of Evidence in Dempster-Shafer Theory." Technical Report, Sandia National Laboratories, SAND 2002-0835 (2002).
41. Squicciarini, A., Bhargav-Spantzel, A., Bertino, E., Czeksis, A.B.: "Auth-SL – A System for the Specification and Enforcement of Quality-Based Authentication Policies." *Proceedings of the 9th International Conference on Information and Communications Security (ICICS)* (2007).
42. Trusted Computing Group: Trusted Platform Module Main Specification `http://www.trustedcomputinggroup.org/` (2003).
43. United States Secret Service and CERT Coordination Center: "Insider Threat Study: Illicit Cyber Activity in the Banking and Finance Sector." `http://www.secretservice.gov/ntac/its_report_040820.pdf` (2004).
44. Wei, Q., Crampton, J., Beznosov, K., Ripeanu, M.: "Authorization Recycling in RBAC Systems." *Proceedings of the 13th ACM Symposium on Access Control Models and Technologies (SACMAT)* (2008).
45. Zhang, X., Nakae, M., Covington, M.J., Sandhu, R.: "A Usage-based Authorization Framework for Collaborative Computing Systems." *Proceedings of the 11th ACM Symposium on Access Control Models and Technologies (SACMAT)*, 180–189 (2006).
46. Zhang, X., Park, J., Parisi-Presicce, F., Sandhu, R.: "A Logical Specification for Usage Control." *Proceedings of the 9th ACM Symposium on Access Control Models and Technologies (SACMAT)* (2004).