CERIAS Tech Report 2007-27

CONTEXT-AWARE ADAPTATION OF ACCESS CONTROL POLICIES FOR CRISIS MANAGEMENT

by Arjmand Samuel, Arif Ghafoor, Elisa Bertino

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

Context-aware Adaptation of Access Control Policies for Crisis Management

Arjmand Samuel*, *Student Member, IEEE,* Arif Gahfoor†, *Fellow, IEEE,* and Elisa Bertino‡, *Fellow, IEEE*

Abstract

Today, most public service delivery mechanisms, such as hospitals, police and fire departments, rely exclusively on digital generation, storage and analysis of vital information. To protect critical digital resources access control mechanisms are employed. The aim is to define rules under which authorized users can access resources required to perform organizational tasks. These rules or *policies* define constraints of time and space on digital resources. Natural or man-made disasters pose a unique challenge, whereby, previously defined constraints may debilitate the ability of the organization to act to its fullest capability. In this paper we propose to employ contextual parameters; specifically, activity context in the form of emergency warnings, to adapt access control policies according to a priori configuration which allows maximum access to critical resources. We also propose an architecture for the detection of crises in the form of activity context and incorporate it in the policy adaptation framework.

I. Introduction

Natural disasters are characterized by floods, tornadoes, volcanic eruptions and earthquakes. Man-made disasters are characterized by wars, industrial accidents and more recently, acts of terror. Management of disaster calls for emergency services such as fire, police and health care to provide services in excess of what they offer otherwise [1]. Similarly, services like banks, insurance and traffic management, are expected to expand their capabilities and provide services

*Purdue University, Electrical and Computer Engineering; West Lafayette, IN 47907. amsamuel@purdue.edu. Corresponding Author

†Purdue University, Electrical and Computer Engineering; West Lafayette, IN 47907. ghafoor@purdue.edu.

‡Purdue University, Department of Computer Science; West Lafayette, IN 47907. bertino@purdue.edu.

in a crisis situation. However, these services rely heavily on on-line access to critical data ranging from rescue coordinates to electronic health records [2]. The onset of crisis demands that the operating envelope in terms of users, duration and place of work be enhanced. In this respect a critical requirement is that data be made accessible to users that may be different from the users who have access to those data under normal conditions.

In most organizations, access to digital resources, such as critical data, is controlled by defining access control rules or *policies* [3], which dictate the conditions or constraints under which authenticated users are granted access privileges. A crisis situation may require that policies governing access under normal situations be quickly adapted to allow timely access to information for crisis management. Detection of a crisis is thus at the heart of adapting access control policies for handling an emergency situation. Government agencies, such as those listed in Table I, warn citizens of the onset of a crisis. Effective and automatic detection and use of such warnings for an organization's access control policy adaptation is the challenge we address in this paper.

In this paper we address such challenge by investigating two key issues:

1) How to detect and comprehend the activity context in case of a crisis or emergency, in a transparent, distributed and flexible manner.
2) How to use activity context to adapt access control policies for continued access to critical resources.

Adaptation of network policies has been addressed by a number of researchers, notable among them are [4], who propose a high level language to formulate policies and adapt network resources based on request parameters. However, in this work the policy itself does not adapt. XACML [5] also takes into account context parameters in access control decisions; however the approach is very generic and thus a lot has to be provided by the underlying application. Also XACML does not support the notion of triggers. To the best of our knowledge, adaptation of access control policies in a crisis situation, based on contextual parameters has not been addressed so far, which is the focuses of this paper.

The remainder of this paper is organized as follows. Section II provides insight into activity context and outlines a categorization based on dependence of activity contextual parameters. Section III describes the access control policy adaptation methodology for temporal and spatial constraints as well as for changes in user authorizations and role hierarchy. Section IV describes

an architecture for implementing access control policy adaptation, based on detection of crisis from distributed sources of activity context. Finally, Section V concludes the paper with comments about avenues of future research.

II. Activity Context for Crisis Management

According to [6], context, in the form of identity, time, location and activity is defined as *"Any information that can be used to characterize the situation of an entity."* Activity context is further defined as *"what is occurring in the situation."* One of the major hindrances in developing a generalized representation of activity context is the fact that it is sensitive to the domain under consideration. For example, a list of activities in the health care environment may include emergency surgery, major surgery, minor surgery, etc. As it may be noted here, this taxonomy is unique to health care and has no relevance to the banking environment. While it is not our aim to develop ontology of the activity context, however, in order to further our understanding of the activity context, we define five types of activities based on their inter-dependence on other contextual parameters as well as dependence on each other.

An activity contextual parameter consists of discrete states, each of which has specific meaning in terms of the domain in question. These states form the parameter values of the activity parameter. An example is the activity context parameter of "Homeland Security Advisory System (HSAS)" [7] which has five states, namely: Low Risk, Guarded Risk, Elevated Risk, High Risk and Severe Risk. While reporting the value of HSAS, one the the above mentioned values will instantiate the HSAS parameter.

Type I Activities are independent of other contextual parameters (time, location, identity), as well as of other activities. Type I activities may have internal and non-obvious links to other activities, but the parameter value does not define any such links. Example of such activities is HSAS, where a Guarded Risk value does not reveal any time, location or user dependence. However, internally, DHS may arrive at a specific value based on other activity parameters not revealed to the public.

Type II Activities depend on other activities for its parameter states. The activities on which type II activities depend may be independent or may depend on other contextual parameters and activities.

Type III Activities depend on time contextual parameter. With each state of this activity, time

information is also presented. An example is FEMAs heat warning, which may be in-effect at certain times of the day.

Type IV Activities depend on location contextual parameter. The relevant location information is also passed along with its value. An example is the NWS Severe Thunderstorm warning, which is relevant only if location information (county, city etc) is associated with it.

Type V Activities relate to a specific set of users only. Users not in this set are not affected by type V activities. An example is the Nuclear Power Plant Emergency by FEMA. While this emergency may prompt Radiation emergency by CDC for the larger community, it is relevant to the users working at the power plant only.

Responsible Agency	Activity Parameter	Possible Values (not exhaustive)	Significance
Department of Homeland Security (DHS)	Homeland Security Advisory System (HSAS)	Low Risk, Guarded Risk, Elevated Risk, High Risk,Severe Risk	Threat level to protected resources
National Weather Service (NWS)	Severe Thunderstorm Warning, Flash Flood Warning, Flood Warning, Tornado Watch, Severe Thunderstorm Watch, Red Flag Warning, Snow Advisory, High Surf Advisory, etc	Parameterization of these alerts not available. However detailed advisories are released	Weather warnings closely linked to the geographic location context as well as time
Center for Disease Control (CDC)	Bio-terrorism, Chemical, Mass Trauma, Natural Disasters and Severe Weather, Radiation, etc	Parameterization of these alerts not available. However detailed advisories are released	Closely linked to time and location context parameters
Federal Emergency Management Agency (FEMA)	Chemical, Dam Failure, Earthquake, Fire or Wildfire, Flood, Hazardous Material, Heat, Hurricane, etc	Parameterization of these alerts not available. However detailed advisories are released	Closely linked to time and location context parameters

TABLE I

PROMINENT ACTIVITY CONTEXT PARAMETERS SIGNIFYING A PUBLIC CRISIS

It may be noted here that most real-life activities may be a combination of the above defined five activities. Take the example of Flash Flood Warning by NWS which includes location

(county, city, etc) and time and may also depend on Severe Thunder storm and Gale warning. In Table I, we list prominent activity context parameters released by government and public entities to signify some form of disaster or crisis. Note that among the four quoted government agencies, there is a significant overlap of the type of alerts that they generate. As an example CDC reports Radiation as one of the alert parameters, which according to FEMA may be covered by Hazardous Material alert, Nuclear Power Plant Emergency and Terrorism alert. The same alert may be abstracted by DHS in the form of a raised HSAS. Crisis Activity parameters generated by multiple agencies provide users a unique method of ascertaining the correctness of these alerts. Moreover those parameters often contain relevant data, such as data on space and time, that combined, provide a rich set of context information. Such information can be used to drive the adaptation of the access control policies.

III. Access Control Policy Adaptation for Crisis Management

Access control policies are typically designed by policy administrators for day-to-day operations of an organization. Those policies often include temporal and spatial context parameters, so that access is allowed or denied depending on the values of these parameters when access is required. However, in case of a crisis, protection requirements for critical information sources shift in favor of more readily accessible information to authenticated users at possibly different times and locations. Specifically the temporal and spatial constraints need to be redefined for catering to the changed situation (see Figure 1(a)). It is important to notice that access control policy adaptation may take place along different dimensions, and it is thus necessary to understand which are the adaptation dimensions relevant for emergency management. Based on the previous discussion we have a list of requirements for access control policy adaptation in a crisis environment. Such requirements, listed below, represent the relevant adaptation dimensions of our approach.

- Changes in temporal constraints in terms of the time of day at which a certain task can be performed and the duration of the task.
- Changes in spatial constraints to allow users to access information in emergency locations, in addition to the normal operating environments.
- Enforcement of enhanced or relaxed identification and authentication requirements of users authorized to access digital resources.

- Escalation and revocation of privileges.

A key issue for the development of an adaptive access control model, supporting the above adaptation requirements, is the adoption of an access control model able to take context information into account. The natural choice in this respect is the one represented by Role Based Access Control (RBAC) model and its extensions. RBAC has emerged as a de-facto model for specifying security requirements of large organizations. Its strength lies in the definition of user roles more akin to the functional responsibilities of users in the organization and abstracting object permissions as roles [8]. Roles can be organized into hierarchies. Through the hierarchy mechanism one can easily organize the set of roles in different groups, and thus clearly distinguish between roles to be used under normal situations and roles to be used in emergencies. The Generalized Temporal RBAC (GTRBAC) incorporates a set of language constructs for the specification of temporal constraints [9]. While temporal based access control models are well suited for enforcing access control decisions with respect to non-mobile users, they are not designed for handling mobility aspect of users employing mobile computing devices and moving among secure and insecure locations. Generalized Spatio-Temporal RBAC (GST-RBAC) [10] adds spatial sensitivity to GTRBAC and support for rich spatial constraints.

We now describe a context-aware policy adaptation methodology based on semantics outlined in GST-RBAC. Each role in GST-RBAC has a state associated with it; the possible states are *enabled*, *active* and *disabled*. The enabled state indicates that the role can be activated by a user. The active state indicates that the role has been activated by the user. The disabled state indicates that the role cannot be used in any user session. A prominent feature of GST-RBAC is the definition of temporal as well as rich spatial constraints which define the time-based and location-based conditions under which a role can be enabled, activated and disabled. Another interesting feature of the model is the notion of enabled and disabled constraints. A constraint may be enabled for a duration or at a location, as a consequence of a run-time events or triggers. We use this feature to adapt constraints in an access control policy in case of a crisis.

A. Adaptation of Constraints

The constraint adaptation strategy is the first crucial element of our adaptation approach, which is based on the use of GST-RBAC. To this end, we introduce a categorization of constraints, both

temporal as well as spatial, into *"Normal Constraints (NC)"* and *"Crisis Constraints (CC)"*.

"Normal Constraints (NC)" are constraints, both temporal and spatial, defined for day-to-day operations of the organization. NCs are enabled in the policy by default, and are applied to all requests being generated by users under normal circumstances.

"Crisis Constraints (CC)" are constraints specifically defined for the operation of the organization in case of a crisis. CCs include temporal and spatial constraints specifically related to a crisis such as working round the clock from makeshift locations. By default the CCs are disabled under normal circumstances and are enabled during a crisis.

The access control policies are composed by policy administrators of the organization. Definition of NCs and CCs is a crucial part of this exercise and is carried out in conjunction with the crisis managers and emergency responders.

Activity context parameters are used to enable CCs and disable NCs at runtime with the help of events/triggers. As discussed in Section II, various types of activity context depend on time, location and other activity context parameters. In this regard, Type III activity event will enable temporal constraints which may be part of CCs and will have no effects on spatial constraints part of CCs. Also, Type III activity events will disable temporal constraints part of NCs and will have no effect on spatial constraints in NCs.

Type IV events will disable spatial constraints part of NCs and will enable spatial constraints of NCs. Further, Type IV will also enable spatial constraints part of CCs and will disable spatial constraints of CCs.

B. Adaptation of User Authorization

In order to effectively cater for Type V activity (activities dependent on users), we introduce two classes of users, namely: *"Weakly Enforced Users (WEU)"* and *"Strongly Enforced Users (SEU)"*.

The *"Weakly Enforced Users (WEU)"* are users who may be in the field and may not require authentication in the event of a crisis. The WEUs may access data, for viewing or update, which may not be critical in nature and/or may not effect the integrity of the over all crisis management process. The need for introducing such user category is based on the recognition that in an emergency situation it may not be always possible to authenticate users who are in

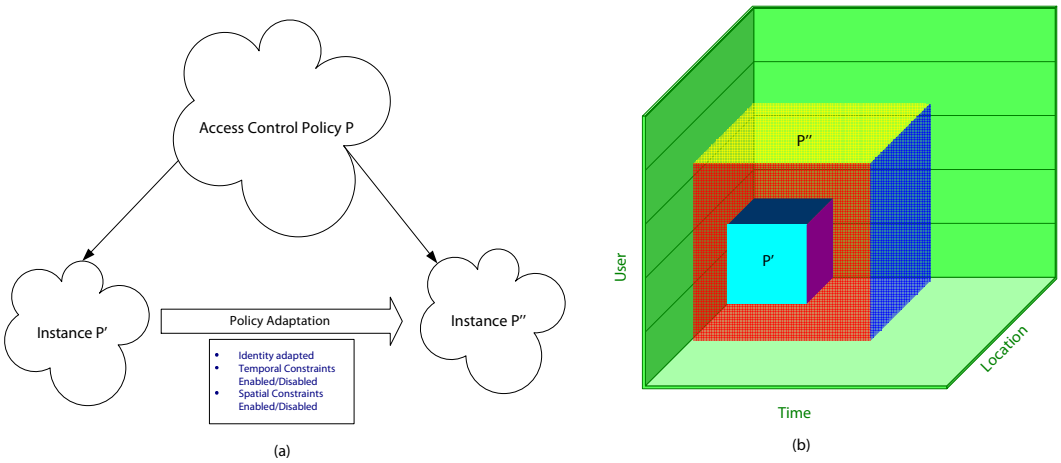

Fig. 1. (a)Adaptation of access control policy in a crisis. Access control policy P is the composed policy containing both directives for use during normal circumstances, as well as those for use in a crisis. P' is the instance of the policy for normal circumstances, and P" is the instance for use in an emergency.(b) Privilege escalation in a crisis; privilege envelope of policy instance P' is escalated to privilege envelope of policy instance P".

the field. In such cases, it may be necessary to trade off strong assurance for data confidentiality and privacy with the need to make those data readily available.

"Strongly Enforced Users (SEU)" are users who may be responsible for critical information access or actions. These personnel could be in the role of a manager or head of surgery etc, and the system may demand them to possess credentials (in terms of user IDs etc) for authentication even in the case of a crisis. Strong authentication is required to protect against attackers attempting to compromise the crisis process management by falsifying information and accessing information sensitive for the crisis management.

The assignment and enabling of WEUs and SEUs in the access control policy is based on Type V activity context in the form of an event or a trigger. It also requires a careful assessment of which resources are highly critical for the crisis management process.

C. Adaptation of Role Hierarchies for Privilege Escalation

RBAC offers role hierarchy as a means for roles to inherit permissions and users from each other. Role r_1 inherits role r_2 (expressed as $r_1 \geq r_2$)if all permissions of r_2 are also permissions of r_1. In this case r_1 acquires all users of r_2. Clearly, defining a hierarchy between roles affords a means of transferring privileges from one role to the other, a feature we use to adapt privileges

in case of a crisis. By defining role hierarchy which is valid only in a crisis, privileges of a role can be handed over to other roles so as to allow maximum number of users the ability to perform certain tasks. Similarly, disabling role hierarchy can also be used to disable subordinate roles. GST-RBAC allows roles to be enabled and disabled as a result of a runtime trigger or event. In the case of a crisis, adapting role hierarchy can take one of the following three forms. Note that the runtime trigger or event in each case is the activity context. The type of activity context will dictate the manner in which it will affect a particular case.

1) Enabling or disabling a role hierarchy in a crisis, results in all subordinate roles (in the hierarchy) to be enabled or disabled, respectively. This mechanism allows escalation or revocation of privileges in a crisis. This can be useful in the case of an unrestricted (both temporally and spatially) role hierarchy. Enabling/disabling of roles based on runtime time triggers and events has been formalized in [9].
2) Adaptation of time based role hierarchy allows roles to inherit privileges at a different time of the day, not allowed previously, and for longer (or shorter) duration. Since the *"enable time"* and *"activation time"* role hierarchies [9] are driven by the respective temporal constraints, their adaptation is achieved by the enabling and disabling the relevant temporal constraints as defined in Section III-A.
3) Adaptation of spatial role hierarchy allows roles to inherit privileges at locations not previously allowed. Enabling/disabling *"simple spatially restricted"* and *"topologically related"* role hierarchy [10] is achieved by the spatial constraint adaptation methodology described in Section III-A.

D. Privilege Escalation and Revocation for Crisis Management

Adapting access control policy in the event of a crisis allows users to perform their tasks at times and places defined specifically for such an eventuality. The operating envelope of the two instances of policy P, before adaptation (P') and after adaptation (P") is depicted in Figure 1(b).

Note that in case of a crisis, authorized users are granted privileges in addition to the ones held by them prior to the declaration of a crisis. For example a physician authorized to work in a hospital from 9 AM to 5 PM in normal circumstances may now be authorized to work for 24 hours, without a break. Similar is the case of location, where a nurse may be authorized to access records of a patient while at the nursing station only; however, in case of an emergency,

the nurse may be allowed to access the same records throughout the out-patient wing of the hospital. Figure 1(b) depicts the privilege envelope expanding in case of a crisis, compared to the smaller one before a crisis.

It is however important to notice that, in the event of a crisis, privileges may also be revoked. A case in point is that of non-essential personnel in a hospital who may not be allowed to work in a crisis so that system load can be reduced. An example is that of a grad student from a neighboring university working on her PhD dissertation by conducting statistical analysis of hospital data to pin point a trend in number of patients committing suicide after consulting a specific physician. While such analysis is not useful in the event of a crisis, it may take up precious system resources.

Using the above mentioned privilege escalation and revocation techniques for GST-RBAC, privileges can be effectively escalated and revoked based on the activity context. The use of events and triggers together with constraints and role hierarchies provide the necessary mechanism to achieve the desirable escalation and revocation of privileges in a crisis.

IV. Access Control Policy Adaptation Deployment Architecture

In this section, we present the system architecture implementing the proposed access control policy adaptation. In particular, we apply the adaptation methodology developed in the preceding sections to a crisis situation in an organization. Although the architecture has been developed with RBAC model in mind, it is general enough to be applied to any context aware access control model.

The architecture is depicted in Figure 2. We now highlight the functionality of key components by describing two scenarios; one with a request under normal circumstances and the other with a request in a crisis situation.

A. Access Control Under Normal Circumstances

The access control policy is composed and stored in the *Access Control Policy Base*. In addition to other components of the policy (as mentioned in [9], [10]), temporal and spatial constraints, classified as NCs and CCs are also stored. The *Policy Instance Generator* creates an instance of the policy with the NC constraints.

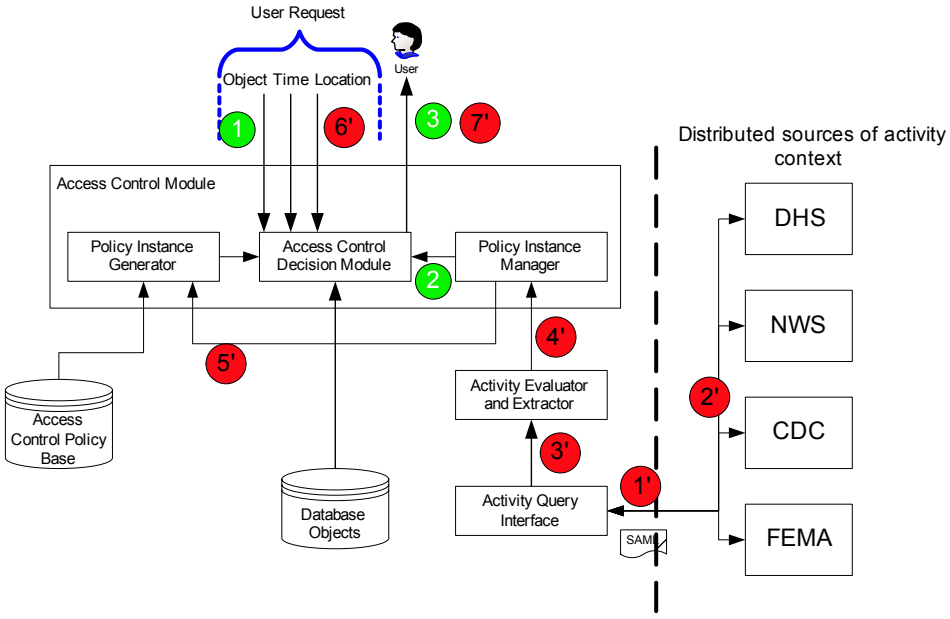

Fig. 2. Access control policy adaptation architecture in a crisis.

Requests from the users in the form of the objects to be accessed, time of request and location of request, reach the *Access Control Decision module (ACDM)*, depicted as label 1 in Figure 2. The decision concerning whether to grant the access is made by the ACDM in conjunction with the *Policy Instance Manager (PIM)* which stores the current instance of the policy. The requested object is retrieved from the *Database Objects* and sent to the user (depicted as label 3 in 2).

B. Access Control for Crisis Management

In order to detect a crisis, the *Activity Query Interface (AQI)* queries the distributed sources of activity context (label 1' in Figure 2). Querying can be based on polling, or AQI can be designed as a passive component waiting for the sources to send out activity context parameters. Either way, the activity context is sent out by the sources (label 2' in Figure 2) and AQI forwards it to the *Activity Evaluator and Extractor (AEE)*. The AEE separates the various types of activity context (as defined in Section II). Note that the types of activity context defines the type of constraint it will affect. As an example, Type III activity context is related to time, and will affect the temporal constraints in the access control policy. AEE is also responsible to correlate

similar contextual parameters received from multiple sources. An example in this case would be the case of NWS providing thunderstorm warning in a particular area along with FEMA releasing a similar alert. In the rare occasion of a mismatch between alerts sent out by different agencies, AEE makes a selection based on predetermined priorities by the policy administrators.

The individual activity context parameter (along with its type) is sent to PIM (label 4' in Figure 2). PIM requests a different policy instance based on the value of the activity context and the CCs in the policy (label 5' in Figure 2). The Policy Instance Generator generates the new instance and loads it in the ACDM (label 6' in Figure 2).

A new request generated by the user (label 7' in Figure 2) is received by the ACDM which grants privileges to the user under a different policy instance. The changed privileges allow the user to access database objects (label 8' in Figure 2) which could not be access under normal circumstances.

V. Conclusion

In this paper, we have presented a comprehensive approach to access control policy adaptation in a crisis situation. The basis of this adaptation is the ability of the access control mechanism to sense crisis and follow a well defined methodology for adaptation. Specifically, we categorize activity context parameters based on their dependence on other contextual parameters and activities. Our proposed policy adaptation methodology describes a step-wise approach to adapting temporal and spatial constraints, user identities and role hierarchy. We also propose an architecture for implementing the adaptation methodology.

A number of issues remain to be addressed. The issue of policy verification during and after adaptation needs to be addressed and a comprehensive verification for the original policy as well as for the adapted policy needs to be devised. The proposed architecture is based on the use of an alert system provided by government agencies, which at present mostly consists of RSS (Really Simple Syndication) [11] feeds. A more machine readable format is required for systems to autonomously query activity context relevant to a crisis using technologies such as Service Oriented Architecture [12].

REFERENCES

[1] C. M. Pearson and J. A. Clair, "Reframing crisis management." *The Academy of Management Review*, vol. 23, no. 1, pp. 59–76, 1998.

[2] W.-T. Wang and S. Belardo, "Strategic integration: A knowledge management approach to crisis management." in *HICSS*, 2005.

[3] S. L. Osborn, R. S. Sandhu, and Q. Munawer, "Configuring role-based access control to enforce mandatory and discretionary access control policies." *ACM Trans. Inf. Syst. Secur.*, vol. 3, no. 2, pp. 85–106, 2000.

[4] L. Lymberopoulos, E. Lupu, and M. Sloman, "An adaptive policy-based framework for network services management." *J. Network Syst. Manage.*, vol. 11, no. 3, 2003.

[5] OASIS. (2007, May) XACML. [Online]. Available: http://www.oasis-open.org/

[6] G. D. Abowd, A. K. Dey, P. J. Brown, N. Davies, M. Smith, and P. Steggles, "Towards a better understanding of context and context-awareness." in *HUC*, 1999, pp. 304–307.

[7] DHS, "Citizen guidance on the homeland security advisory system," 2007, p. http://www.dhs.gov/xlibrary/assets/CitizenGuidanceHSAS2.pdf.

[8] R. S. Sandhu, D. F. Ferraiolo, and D. R. Kuhn, "The NIST model for role-based access control: towards a unified standard." in *ACM Workshop on Role-Based Access Control*, 2000, pp. 47–63.

[9] J. Joshi, E. Bertino, U. Latif, and A. Ghafoor, "A Generalized Temporal Role-Based Access Control Model." *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 1, pp. 4–23, 2005.

[10] A. Samuel, A. Ghafoor, and E. Bertino, "Framework for specification and verification of generalized spatio-temporal role based access control model," *CERIAS Technical Report*, vol. TR 2007-08, 2007.

[11] D. Chmielewski and G. Hu, "A distributed platform for archiving and retrieving rss feeds," *icis*, vol. 0, pp. 215–220, 2005.

[12] K.-P. Eckert, "The fundamentals of web services." in *The Industrial Information Technology Handbook*, 2005, pp. 1–15.