Content-sensitive, Temporally Adaptive Metadata

Brendan J. Gilbert[1], Raj Sharman[1], Manish Gupta[1], H.R. Rao[1], Shambhu Upadhyaya[1]
Kenneth P. Mortensen, Esq. [2]
[1] The State University of New York, Buffalo
{bg1,rsharman,mgupta3,mgmtrao,shambhu}@buffalo.edu
[2] Privacy Office, U.S. Department of Homeland Security; U.S. Department of Justice
kenneth.mortensen@dhs.gov

Abstract–**Role-based access is the most commonly used method for providing access to information systems. Roles are secured through design principles such as least privilege and separation of duties. However, during emergency situations, system availability to first-responders and emergency coordinators through privilege escalation has proved to offer tremendous benefits. While need for privilege escalation had received much attention, little research and focus has been given to area of ensuring security of information after the emergency. Focus of the paper is secure return of access privilege levels to normalcy after the emergency situation and resulting risks. This paper discusses some models for managed privilege escalation, using a deterministic finite state machine as a framework to select sets of context-sensitive and temporally adaptive metadata, with environmental and temporal state transitions. The framework is demonstrated through its application to a historical scenario whose result could have been improved by having such a framework in place. Risk assessment discussions are also provided to ensure that reliable and secure roles are designed (for emergency) and secure transitions occur (during and after emergency).**

I. Introduction

The Health Insurance Portability and Accountability Act (HIPAA) protects the disclosure of Protected Health Information (PHI) by limiting access. This has resulted in databases containing PHI to enact different levels of user privilege to provide the minimum amount of information necessary for the intended use [1]. During emergency situations, however, PHI's availability to first-responders through privilege escalation saves lives, as evidenced by the disclosure of PHI to locate tuberculosis patients evacuated across the United States in the response to Hurricane Katrina [2]. The importance of the availability of PHI during duress is supplemented by the Department of Health and Human Services' development of an emergency responder electronic health record to assist with assembling medical histories to support emergency relief efforts and develop an interoperable record of PHI that can be disseminated quickly as needed [3]. A concern of privilege escalation is that once data is made available at a certain privilege, it cannot be made unavailable with assurance.

However, the impact of availability can be assessed and mitigated by creating structured systems with set access privileges. This article addresses some possible models for managed privilege escalation, using a deterministic finite state machine as a framework to select sets of context-sensitive and temporally adaptive metadata, with environmental and temporal state transitions. The framework is demonstrated through its application to a historical scenario whose result could have been improved by having such a framework in place.

II. Previous Work

Privacy in healthcare is of a high concern to our society. HIPAA has set expectations and made breaches of this privacy more actionable. Recently this concern has been punctuated by events such as the theft of 40,000 patient records containing the names, phone numbers and social security numbers on April 11th 2008. The scope of the theft at New York-Presbyterian Hospital/Weill Cornell Medical Center in Manhattan was uncovered by a federal investigation and an internal audit, the hospital said [4]. The exposure of this information could cause considerable damage to the patients, whose information has been stolen and can be misused in multiple ways, including perpetrating financial frauds. However, that same identifying data can also be used during emergency situations to save lives. In the wake of Hurricane Katrina's devastation, the U.S. Department of Health and Human Services' Office for Civil Rights issued a memorandum affirming that the Privacy Rule "…allows patient information to be shared to assist in disaster relief efforts, and to assist patients in receiving the care they need." [4] In response to a congressional mandate in the Health Insurance Portability and Accountability Act of 1996 (HIPAA), HHS issued regulations entitled *Standards for Privacy of Individually Identifiable Health Information*. For most covered entities, compliance with these regulations, known as the Privacy Rule, was required as of April 14, 2003 [15]. For covered entities using or disclosing PHI, the Privacy Rule establishes a range of health-information privacy requirements and standards that attempt to balance individual privacy interests with the community need to use such data.

A concern unaddressed by the Privacy Rule's permissions to distribute PHI is the effect of the memory of responders. Once privileged data has been made available to a user, it cannot assuredly be made unavailable. Although access can be rescinded, the ability to recall PHI is not removed with removing access due to memory. From this, it is clear that in order to mitigate the likelihood of PHI being used for unethical purposes, even during emergencies the best practice of least-privilege access should be adhered to. The principle of least privilege has been described as important for meeting integrity objectives [16]. The principle of least privilege requires that a user be given no more privilege than necessary to perform a job. Ensuring least privilege requires identifying what the user's job is, determining the minimum set of privileges required to perform that job, and restricting the user to a domain with those privileges and nothing more. By denying to subjects transactions that are not necessary for the performance of their duties, those denied privileges couldn't be used to circumvent the organizational security policy [6].

Previous work examining context-sensitive access in the scope of security has provided a basis to shift away from the Role Based Access Control (RBAC) model due to the fact that the model does not accommodate flexibility [5]. Roles are strictly defined as well as the access rights provided to them. Users of the system fall into one or more of these roles, and therefore have the associated access rights due to membership. However, this model lacks any awareness or notion of contextual factors as a determinant of access privileges for a given role, nor a sense of a progression through states over time.

A model permitting roles based on context has augmented this, providing a Generalized Role Based Access Control (GRBAC) model. One example of GRBAC is the defining of different user roles based on day of the week. For instance, a payroll administrator in GRBAC system may be allowed to make payroll modifications only on a certain day of the week and resources may be accessible during certain hours of the day [6]. A second example given is to restrict access to services offered by a transportation company to legitimate users of their service—for example, being able to use WLAN services provided to travelers on a railroad carriage. In this case, the environmental role is restricted to users whose location moves at the same speed as a GPS locator on the train, to ensure that they are legitimate customers of the railway [7]. The GRBAC model allows for changes based on environmental context, but does not consider time or the different access privileges of roles based on different contextual triggers.

Applying environment roles that include a shift in the required escalation of privileges due to disaster response, and then de-escalate the privileges incrementally as the situation comes under control or as time passes, is absent from existing work. Particularly with information as sensitive as PHI, the ability to plan for changes to access controls is required to be able to assess the amount of risk associated with privilege escalation, required to engineer a business continuity plan [8, 9].

III. Contextuality and Temporality

How can least-privilege be maintained during emergency situations? A set of metadata sets, containing the environment roles' access control lists of different parts of the dataset, would yield adaptability to various contexts as well as provide a basis for an incremental, staged return to the least-privilege state over time. This can be viewed as a deterministic finite state machine, with the states representing various sets of roles that provide levels of privilege to responders. Contextual inputs cause transitions from the least-privileged state of metadata to states of progressively higher privilege, depending on the magnitude of the context. In states of escalated privilege, the passage of time or the occurrence of events to rectify the situation that requires higher privilege cause transitions to states of less privilege. This interpretation gives rise to a linear model of the state machine, as shown in Figure 1. .

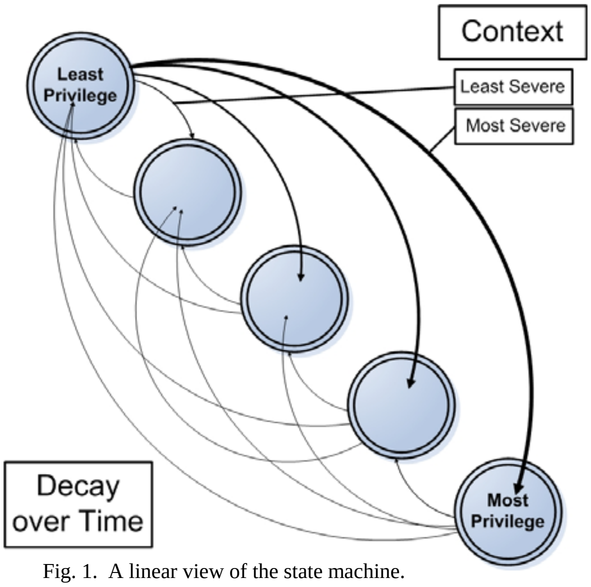

Fig. 1. A linear view of the state machine.

With these models defined, we will examine how having a contingency plan define these states and the privileges associated with each role defined by them improves the integrity of sensitive data, in the event that privilege escalation is required.

The aim of the model's defining levels of privilege is to facilitate the estimation of risk that may arise from disclosure (confidentiality) and un-availability of information. During an emergency situation necessitating the release of PHI, it is difficult to gauge the scope of the distribution. For example, during the relief efforts for Hurricane Katrina, PHI was released to any organization providing support, such as the American Red Cross. However, since these organizations are not covered entities within the scope of HIPAA, it is likely that the sensitive PHI of some inhabitants of New Orleans are now available to many individuals from organizations that cannot be held liable for the compromised privacy of those whose PHI was distributed. The HIPAA Privacy Rule does not extend to organizations that are provided access to PHI to aid in their contribution to the relief effort [5].

Although access control cannot prevent this situation from happening again, the scope to which the released PHI is made available can be estimated through risk assessment. By planning for the amount of privileges extended to various groups of users with a state machine as shown above, the amount of risk from the scope of availability of sensitive data can be gauged and accounted for. In an effort to maximize the ability to respond to breaches of confidentiality, levels of logging can also be specified to mitigate the risk associated with the increased availability of sensitive data in an escalated privilege state.

IV. A SCENARIO

Treatment for tuberculosis requires at least six months of supervised medication. When Hurricane Katrina struck New Orleans on August 29, 2005, there were 130 residents undergoing this treatment. Ensuring these patients remained on regimen and had an adequate supply of medication took a high priority during the evacuation, to prevent epidemics breaking out over the several states evacuees sought refuge in.

The CDC's controlled efforts of this situation were very successful, getting all 130 patients back into the pharmacological fold by October 13, 2005. PHI was used in order to locate the difficult to find patients, spread over states as distant as Washington and Massachusetts. HIPAA's regulations required the creation of limited arrangements with pharmacies to cross-reference prescriptions dispensed to the tuberculosis patients' information. The CDC admits, "Prearranged agreements of this type, applicable to various health-related emergencies, would have facilitated faster location of patients," as well as "standardize[d] electronic health records" and "HIPAA-compliant platforms for sharing information." [2]

V. PRACTICAL IMPLICATIONS

Risk associated with any state of metadata, governing the roles of users and the privileges granted to each role, is a function of 1) the number of people with access, 2) the amount of access provided to those people, and 3) the amount and sensitivity of PHI that is distributed. This amount of risk is mitigated by the granularity of the result sets: for example, instead of providing a user all of the emergency-contact individuals for a certain missing tuberculosis patient, the user could only see those who were living in New Orleans and therefore might be able to provide input as to their current whereabouts. Additionally, the user would not be able to see the relationship between that contact person and the individual. The restrictions placed on result sets could dynamically change based on contextual triggers modifying the needs of the database's users. The triggers would be initiated during emergencies where adaptive roles would be re-assigned to new (or additional) set of users who would need access to information (per new roles' privileges) to respond to emergency. Possible amounts of risk associated with these are documented in Table 1.

TABLE 1
GRANULARITY OF QUERY STATES, AND ASSOCIATED RISKS

	Access to PHI		**Possible Effects**
High Risk States	PHI in excess of required amount	Excess PHI, too small result set	Liability for identity theft Loss of public trust/goodwill
	Very little PHI, no range required	More PHI, must define search range	Mitigated chances of identity fraud
Low Risk States	Very little PHI, must define search range		Hacking authorized account
	Denial of any PHI access		Server exploits to modify privileges

Another possible mitigating factor is the role of trust, which would reduce the liability of individual users. Although it is probably safer to assess risk as a matter of worst-case scenarios and thereby discounting the influence of trust, if trust is a significant differentiating factor between users it should also be included. Providing centralized, trusted users a higher level of access than their peers may achieve the same level of operating efficiency while reducing risk based on the number of people with access. For example, providing one "trusted" supervisor, who is a full-time employee of the relief organization with higher access than his or her volunteers might enable the same level of efficiency of locating missing patients due to directive management, but the risk of fraud would be lower, due to fewer individuals having the amount of access that the supervisor does.

In addition, there should exist certain restricted states or duties for which privilege should not exist to users who are outside the typical assignees of that privilege. These states would have no transitions that lead into them, and may have either no transition to lead out of them or an alternate progression of state transitions than normal users to ensure that they are kept separate. Possible distributions of risk and restricted states are visualized in Figure 2—note that each state in the machine would have an associated risk level (R_N) that may provide an incrementally higher amount of risk than the previous level (the case in the concentric diagram on the right) or may provide similar amounts of risk but not encompass all of the risks of the previous state (the case in the irregular diagram on the left). These amounts of risk are highly variable dependant on the system being evaluated, and are the domain of the system or database architect to consider.

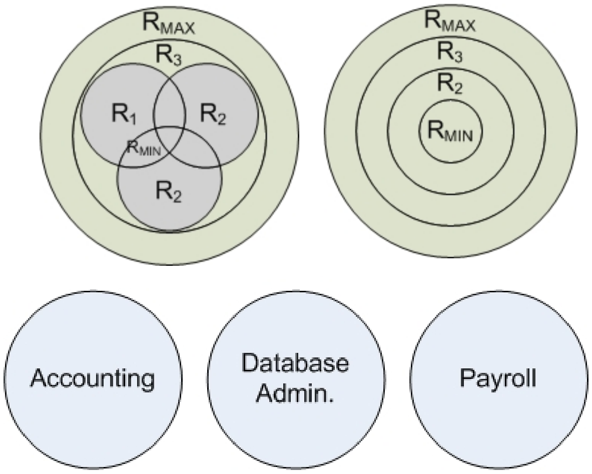

Fig. 2. Possible distributions of risk ($R_{MAGNITUDE}$) and restricted states/duties.

VI. Applied Examples

Applying this model to the after-effects of the September 11th, 2001 attacks, we would see a progression to a very high level of privilege offered to those acting on behalf of national security, with relaxation over the passage of time back to a higher default lowest-privilege level. The steps of returning to the lowest-privilege state across the top of the matrix model are in this case omitted to reflect the higher amounts of access privileges available to all users responsible for the reduced privacy in effect post-attacks. Making such a change would require that new orders be determined for risk progression, based on the new lowest-privilege state. Progressions to higher privilege states could be based on the National Threat Advisory provided by the Department of Homeland Security.

The controlled disclosure of PHI could also be used to ensure that the homes of Hospice patients are given a higher priority for power restoration during a blackout. Following an unexpected blizzard in October of 2006, a million Western New York residents lost power for at least a day, with 350,000 households losing power for the majority of a week. The full death toll for the storm included 13 people. Due to triage the Army Corps of Engineers, as part of the relief effort, supported churches and missions first due to the number of people seeking refuge there. Residents who needed electricity to power home medical equipment "were in dire straits," according to the Corps [10]. If the Corps had access to disclosed PHI, including addresses of local residents requiring power to use home medical devices—for example, a dialysis machine, or a respirator—the residents could have been prioritized in the power restoration process, lowering the number of casualties.

VII. Discussions on Risk Management

While system access structure should be adaptive to accommodate privilege escalation during emergencies, there are several risk factors that should be accounted for while the system roles and escalation workflows are architected.

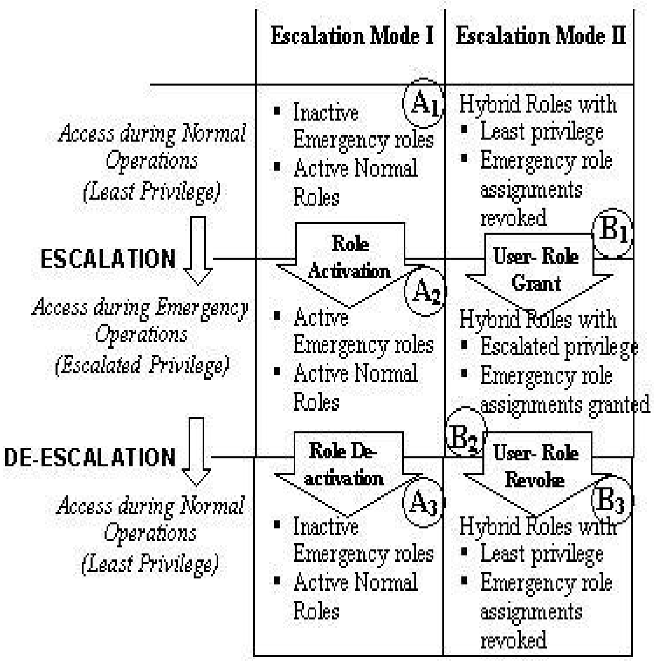

Typically, there are two common methods of granting escalated system access (See Figure 3): 1) ESCALATION MODE I (ROLE – ACTIVATE/DE-ACTIVATE), which involves creating roles with higher authority, according to requirements to respond to a crisis. In this mode, roles remain dormant until a situation of emergency, and 2) Escalation Mode II (User – Grant / Revoke) where the system already has roles that have sufficient privileges to accommodate efficient response to an emergency. In this case, designated users are assigned to relevant roles to perform duties during an emergency. An automated or defined process should be in place to initiate the assignment with least latency, while ensuring information assurance.

In either of the above-mentioned escalation modes, there are several threats that can arise:

a. How to engineer roles for system access that can be used during both - normal operations and emergency situations, while ensuring lowest risks and threats to information assurance.
b. How to make sure that escalations happen without compromising confidentiality, integrity and availability of information contained in the system.
c. What is the more suitable design between 2 escalation modes given the system states and ensuring availability of data.
d. How to incorporate accountability and audit-ability of roles usage during escalation and during emergencies.
e. How to create workflows/processes for efficient de-escalation of privileges, in either of modes, without affecting response to crises.
f. General security questions such as:
    i. What's the impact if an attacker, during emergency, can manipulate the escalation process or system itself to read the system data? What happens if access is denied to the system during emergency?

To aid in asking these kinds of pointed questions, we argue for the use of threat categories by adapting and extending the STRIDE threat model. Developed by Microsoft, STRIDE is an acronym derived from the following six threat categories: Spoofing identity (S), Tampering with data (T), Repudiation (R), Information disclosure (I), Denial of service (D) and Elevation of privilege (E) [12]. In fact, the above threat categories may not be mutually exclusive. A threat that is exploited can lead to other threats. Some threat types can interrelate. It's not uncommon for information disclosure threats to lead to spoofing threats if the user's credentials are not secured. And, of course, elevation of privilege threats are, by far, the worst threats—if someone can become an administrator or can get to the root on the target computer, every other threat category becomes a reality [14]. Conversely, spoofing threats might lead to a situation where escalation is no longer needed for an attacker to achieve his goal. For example, using SMTP spoofing, an attacker could send an e-mail purporting to be from the CEO and instructing the workforce to take a day off. To capture the various nuances of this transaction from the threat focus, and to get a better understanding of the components involved, a workflow representation is developed [13].

Similar applications have been suggested in research for developing a framework for the measurement of security levels of any EBPP [14] system to help security personnel to ensure a higher level of understanding of information assurance issues and proactively engage in elevating security measures and fraud protection in their organizations. We studied the 7 steps risk assessment framework [14] and believe that it can be adapted for managing risks for adaptive role systems. Figure 3 shows two escalation modes with system states before, during and after an emergency represented as A_1, A_2, A_3, B_1, B_2 and B_3. There are various threats, that can arise due to design of a role system for emergency roles, which should be evaluated for risk management. Proper analysis of state levels X_i will ensure that role design is secure and efficient transitions of states take place during an emergency and during restoration. The figure shows 3 states for the system with escalation mode I, where specific and exclusive roles are designed for use during emergency. These roles are only activated during emergency and de-activated after. System states, using escalation mode II, rely on assigning additional users to existing system roles.

VIII. Future Research and Conclusions

This work is part of an ongoing larger project to examine how to integrate context-sensitive metadata governing privilege escalation into continuity planning. Until now GRBAC has been viewed as a way of moderating user privileges during normal operations—using it to aid in disaster recovery planning, and to understand the risks during plan implementation is an avenue not yet explored, which should be in greater depth.

Despite the crisis of an emergency, proper emergency management should aim to provide a return to normalcy as soon as possible. The temptation to "open the floodgates" of information by relaxing lowest-privilege access to data is strong, particularly if there exists no framework in place to be able to gauge the effectiveness of progressive amounts of requirement reduction. Application of this framework provides a rational way to provide the amount of access needed to save lives, yet not too much access so as to increase the grief of survivors of an emergency through the possible fraud due to information disclosure can bring.

To efficiently implement a series of metadata that fulfill this concept, it would be best to modularize parts of the metadata. A separate relational database of different metadata states could accomplish this with a minimum of redundancy.

A possible extension of this project, especially given the modularized groups of metadata, is developing a regression analysis to forecast the amount of risk that a given state would produce. A regression analysis equation would support the dynamic creation of new states based on a certain maximum acceptable threshold for risk, and also certain rules governing how privileges and roles should be combined to make a logical and efficient system. For example, using the regression coefficients provided from the analysis, an algorithm could dynamically determine the most optimal state to provide as much access as possible within a certain allowed maximal amount of risk, and then trigger a change into that optimal state. As the requirements change, the algorithm could again calculate the most efficient state of metadata to use and then transition into the new, dynamic state automatically.

REFERENCES

[1]. "Uses and Disclosure of Protected Health Information: General Rules." Code of Federal Regulations Title 45, Pt. 164.502(b)(1), 2006 ed.

[2]. Center for Disease Control. "Tuberculosis Control Activities After Hurricane Katrina," *Morbidity and Mortality Weekly Report*, vol. 55, pp. 332-335, 2006.

[3]. United States Department of Health and Human Services. "Emergency Responder Electronic Health Record Detailed Use Case.," Loonsk, John W. Washington, DC: U.S. Department of Health and Human Services, 2006.

[4]. United States Department of Health and Human Services, "Hurricane Katrina Bulletin: Disclosing PHI in Emergency Situations," *Department of Health and Human Services,* 2005. http://www.hhs.gov/ocr/hipaa/KATRIN AnHIPAA.pdf, 20 Dec 2006].

[5]. D. Ferraiolo, and R. Kuhn, "Role-Based Access Control." In *Proceedings of the 15th National Computer Security Conference,* October 1992.

[6]. M.J. Covington, W. Long, S. Srinivasan, A.K. Dey, M. Ahamad, and G. Abowd, "Securing Context-Aware Applications Using Environment Roles," I*n Proceedings of the Sixth ACM Symposium on Access Control Models and Technologies,* Chantilly, VA, 2001.

[7]. B. Hulsebosch, A. Salden, and M. Bargh, "Context-Based Service Access for Train Travelers," *In Proceedings of the 2nd European Symposium on Ambient Intelligence (EUSAI),* Markopoulos et al. (Eds.), LNCS 3295, pp. 84-87, Eindhoven, The Netherlands, 2004.

[8].B.J. Dooley, "Preparing a Business Continuity Plan." *Faulkner Information Services,* 2006.

[9]. "Business Continuity Planning Booklet," *Federal Financial Institutions Examination Council,* 2003.

[10].U.S. Army Corps of Engineers, Buffalo District. Home Page. U.S. Army Corps of Engineers., 2006. http://www.lrb.usace.army.mil/

[9] M. Howard, and D. LeBlanc, "The STRIDE Threat Model". *Writing Secure Code*, 2002 ed. Microsoft Press. Chapter 2: Designing Secure Systems, pp. 38 – 60.

[13] Y.I. Song., H.R. Rao., and S. Upadhyaya, S. "Information Assurance Issues of the Workflow management Systems in E-Banking: An investigation on the modal points with high risk," University at Buffalo, Working paper, June 2003.

[14] G. Tanna, M. Gupta, H.R. Rao, and S. Upadhyaya, "Information Assurance metric development framework for electronic bill presentment and payment systems using transaction and workflow analysis," *Decision Support Systems Journal*, vol. 41(1): pp. 242-261, 2004/2005.

[15] "Clinical Research and the HIPAA Privacy Rule". Retrieved April 19, 2008 from NIH website at http://privacyruleandresearch.nih.gov/clin_research.asp

[16] "Integrity in Automated Information Systems". *National Computer Security, Center*, September 1991.