Using Budget-Based Access Control to Manage Operational Risks Caused by Insiders

Debin Liu*, L. Jean Camp, XiaoFeng Wang and Lusha Wang
School of Informatics and Computing, Indiana University

Abstract

The insider threat has been framed as protection of the network from insiders whose threat level may be unknown to the organization. In this paper, we propose a Budget-Based Access Control Model to mitigate the insider threat. We provide an order of magnitude price for every access right and assign each individual user a risk budget. The price for access is then personalized based on the observed historical behavior of the user. The risk budget represents the amount of risks an organization can tolerate from that employee. Each access right of a user may cost him certain risk points. The incentives come in the forms of punishments and rewards. The punishments are triggered by the risk budget exhaustion. On the other hand, those whose risk behavior is aligned with the organization's risk preferences will be rewarded. The human-subject experimental results demonstrate our model's positive influence on the users' risk behavior. In addition, this work is distinguished from previous risk-based access controls by our modeling of users behaviors, prevention of risk point hoarding and provision of explicit pricing. All risk-based access inherently constrains behavior incentives.

Keywords: Insider Threat, Access Control, Risk Management, Incentive Engineering, Human-Subject Experiment

1 Introduction

Risk-based access control systems [1][2][3][4] argue that any access control system is an attempt to model the organizational risk: the more fine grained the access control systems become, the tighter we bind on the organizational risk. Thus risk-based control systems are primarily motivated by the need to bypass controls. However, the ability to bypass controls is a main cause of internal fraud, which has been proven to be a serious threat. Consider the following examples.

- More than 10,000 patients' personal information may have been leaked by an employee at Johns Hopkins Hospital in an identity fraud scam in 2009. The hospital stated that the patient registration database contained no medical records, but it did contain sensitive data, such as addresses and Social Security Numbers. The hospital emphasized that the breach was not a hacking incident, but that the employee had access to the records as part of her job [5]. In this incident, the employee had no incentive to protect the data.
- The biggest trading fraud in banking history was discovered in 2008. An employee at the French bank Societe Generale, Kerviel, was rewarded with his work at the back and middle offices for a promotion to work in the front office as a junior trader. A year later, in 2006, he had begun creating fictitious trades. These fake transactions were relatively small in the beginning but increased in frequency and in size throughout 2007. By 2008, he had taken massive fraudulent directional positions hidden behind those faked transactions. As a trusted insider, he used his knowledge to circumvent the controls that were in place to avoid detection. By January 2008 when his unauthorized trading activities were finally uncovered, Kerviel had cost the institution $7.2 billion in fraudulent trades [6]. The individual controls did not limit the aggregate risk in this incident.

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 1, number: 1, pp. 29-45
*Corressponding Author, Email: *deliu@umail.iu.edu*

These incidents were committed by *insiders*: individuals who were authorized to use the information system they employed to create enormous risk for their employers. An insider's privileged position gives him the opportunity to easily abuse the trust of his organization. This creates a grave risk to many institutions [7]. Note that these attacks can't be eliminated by straightforward application of risk-based access control. In the first example, the employee at Johns Hopkins Hospital can always hoard risk tokens in order to access and leak patients' sensitive information.

The insider threat has been shown to be extremely difficult to mitigate: compared with an external attacker, an insider can have an intimate knowledge of the security controls in place. Take the insider case at Societe Generale as an example: Kerviel was a trusted insider. His previous work in the back and middle offices provided the knowledge to circumvent monitoring and bypass controls that were in place in order to avoid detection. In fact, Kerviel successfully avoided the bank's internal controls and escaped detection from 2006 to 2007.

Kerviel was later accused of exceeding his authority to engage in unauthorized trades. Banking security professionals believe that a strong control mechanism can protect institutions from similar catastrophic frauds [8]. In our paper, we treat the violation of "exceeding authority" as a signal requiring better risk management and the violation of "unauthorized trades" as a signal requiring an immediate improvement in the access control system. In other words, we view the insider threat as a security problem to risk management and access control systems. By modeling users' risk behavior and using personalized incentives, our model can mitigate the grave risk without unduly constraining employee behavior.

Given that the insiders are usually rational and motivated by realizing their personal gains, we believe incentive engineering can help us manage the organizational risk and mitigate the insider threat. We have illustrated the theoretical potential of incentive engineering in the general case of an insider [9]. Conceptually, we give each user a budget of risk points, and roughly quantify the risk associated with every move the user takes. The incentives are created in the forms of punishments and rewards. If the user exhausts his risk budget before having his employment tasks completed, he would be subject to penalties from his organization. On the other hand, if he avoids risk and finishes his task without exhausting his budget, the organization can reward him. This observation is the foundation of our Budget-Based Access Control Model which manages access and limits the risk of both malicious and inadvertent insider activities.

The use of risk points has been proposed by several risk-based access control systems [1][2][3][4]. All of them use risk tokens as the main enforcement mechanism. However, risk point hoarding is a particular problem for the insider threat. The first contribution our model made is to prevent risk tokens hoarding with the use of the risk budget which is an individual upper limit.

The second contribution our model made is the modeling of subjects' risk behavior. We learn the subjects' behavioral history and mark each of them with an attribute of riskiness. The riskiness describes a subject's risk attitude and risk behavior. In our experimental evaluations, the higher the value of riskiness is the more risk seeking behavior the user has.

The third contribution is the personalized incentives our model provides. The incentives in the forms of punishments and rewards are measured by the consumption of risk budget. The personalized incentives are offered to an insider and make him to behave according to the risk posture set by the organization.

Our model is the only risk-based access control, to our knowledge, that is evaluated using human-subject experiments. The experimental results demonstrated our model's positive influence on the users' risk behavior.

The importance of each of contributions is detailed on insider threat and risk-based access control in Section 2. In Section 3, we describe our Budget-Based Access Control Model based on a real-world access control system. We move on to describe the human-subject experiments and evaluate our design in Section 4 and Section 5. We conclude the paper and describe future work in Section 6.

2 Related Work

An insider is defined as someone within an organization or with access to critical information aspects of the organization [10]. Previous work [11] by the United States Secret Service and Carnegie Mellon University illustrates that the prevalence of computer crimes perpetrated by insiders in US organizations.

Detection and monitoring techniques are widely used to mitigate the insider threat. For instance, [12] describes an approach to detect insiders' misbehavior by monitoring system calls. These techniques usually rely on the knowledge of the insiders' behavior model and a pattern of the malicious actions the insider has to take to achieve his objective. Examples of the model can be found in [13] [14] [15]. These approaches are meant to analyze external threats and are therefore insufficient for mitigating the insider problem. For example, an insider can stay within his assigned privileges and still attack the system abusing the trust inherent in his access, which makes it hard to accurately detect his malicious activities.

Given the insiders are usually rational and motivated by financial gain, we believe that incentive mechanisms can deter them from implementing attacks which, from the institution's perspective, creates organizational risk. The incentives have to be aligned with the interests of the organization and the individual [16] [17]. For example, security incentives that prevent users from performing critical tasks will be subverted or disabled. The core research challenge our design addresses is how to add incentives into access control systems so that either the attacking behaviors incur some cost, or to make the security costs of a misbehaving account visible to the user. Essentially the research question is how to encourage users not to be risk-seeking by utilizing incentives. Consider the need for incentives in the case of Johns Hopkins Hospital.

In this paper, we propose to mitigate the insider threat by combining a risk budget mechanism with access control system. This offers incentives to an insider to behave according to the risk posture set by the organization. In comparison, [1] proposed an access control system that used a market to distribute access tokens where the price may be set by the data owner. The response is static and the system does not evaluate the responses in order to identify the risk posture of user. Horizontal Integration [2] proposes the use of risk tokens and risk calculations to manage access control. Tokens are distributed to employees in a hierarchical approach, by the organization. Again, employees trade tokens for access. FuzzyMLS [3] considers access control as an exercise in risk management. FuzzyMLS also computes a quantified estimate of risk associated with a human subject, related to our conception of risk posture. FuzzyMLS utilizes risk tokens in a zone of uncertainty, a fuzzy or gray area, between permission and denial and proposes an unspecified market for risk exchanges. [4] describes the mechanisms for distributing risk tokens to employees for access control.

All these risk-based access control systems use risk tokens as the main enforcement mechanism. However, none of these risk-based access control systems bound the amount of aggregate organizational risk. They also only consider predefined exceptions, defining a small number of exceptions and maintain a binary view of access control. In addition, the exceptions provided by these risk-based access control systems create vulnerabilities of insider threat. Insiders can always abuse the exceptions to bypass access control.

All access rights in our system are priced in risk points. When a user needs to bypass the control, he has to pay the price from their risk budgets. Exceptions to bypass access controls are then dynamically created through the price-payment interaction. These exceptions are no longer predefined.

Risk point hoarding is a particular problem for the insider threat. Other proposals allowing hoarding explicitly through trading or implicitly through lack of auditing or expiration. Our model prices every access right, evaluates every user's riskiness, and assigns each user a risk budget. A user pays for the exception using his risk points. A user gets punished if he exhausts his budget. In contrast, the fewer risk points consumed, the greater the user's reward. The budget size along with the incentive of punishment bounds the risk caused by each user within the amount of his given risk budget. Therefore, unlike

previous risk-based access control systems, the aggregate organizational risk could be always bounded. Moreover, the incentive in the forms of punishments prohibits insiders' possible unlimited abuses of access rights.

The price generated in risk points reflects the organizational risk caused by the access rights. In general our model provides order of magnitude risks. The price is also determined by the subject's individual risk behavior. This creates personalized incentives to each subject. Users are influenced by the personalized incentives, and thus the payment paid by the users reflects their willingness to obtain the exceptional access. Therefore, we argue the price-payment interactions bridge the gap between organizational risk and the users' willingness, and build an alignment of incentives between the organization and the users. As a result, the incentive in the forms of rewards encourages users to avoid risky activities and unnecessary exceptional access requests.

Our model is also the only access control using risk points and evaluated by human-subject experiments. Our previous work [9] used a scenario of web-browsing. In this paper, we chose stock trading activity to evaluate our model. The experimental results shows that our approach exerts positive influence on rational users' risk attitudes, and reduces the organizational risk caused by the access exceptions.

3 Budget-Based Access Control

An organization is comprised of employees who perform work and complete tasks. An employee is the *subject* of an access control decision. He may be an *insider*. We designed a Budget-Based Access Control Model to mitigate this insider threat. Our model studies each subject's behavioral history, and evaluates his riskiness. Users' riskiness measured to reflect their risk attitude and behavior are updated periodically. We then price every access right in the organization in risk points for each subject. The pricing mechanism considers both the subject's riskiness and the request's risk classification. Those who have a history of risk seeking behavior are charged a higher price. Those who have been risk averse are charged less.

The third mechanism used in our model is the risk budget mechanism, which assigns every subject a set of risk budget parameters based on the subject's organizational position and task requirement. The risk budget parameters include a budget balance, r, a lower risk bound, l, and an upper risk bound, h. When the price of the requested access is below the subject's lower risk bound, it can be accessed for free; when the price is greater than the subject's upper risk bound, the subject can not access it; and when the price is between the lower and upper risk bounds, a price-payment interaction is launched to make access control decision. In this paper, we assume the allocation of risk budgets is done manually by the organization. We will explore the possibility to automate this process in the future.

Risk budget parameters are assigned based on each subject's position and task, which could be conceptually described as the subject's role. This makes it easy to integrate our model to any existing role-based access control. For simplicity of description, here we use the access-control system in a financial institution as an example to illustrate the framework of our model.

Prior work [18] describes an existing real-world access control system within a major European bank. The system, called FUB (Funktionale Berechtigung), is an enterprise-wide role-based access control system. A role is defined as a combination of the official position and job function. Typical official positions could be the ordinary Clerk, Group Manager or Regional Manager. Functions represent the users' daily duties such as a financial analyst, share technician, or internal software engineer. All these data are defined and maintained in a human resources database. A batch job combines such data with FUB every night. Thus, the access control system has a very accurate image of the current organisational status and existing roles. Within FUB the data delivered by the human resources database are linked to applications. When a user starts an application the FUB delivers the security profile that tells the

application the appropriate access rights. Several applications can be accessed by a single role. Each application involves a set of access rights according to the user it delegates.

3.1 *Overview and Framework*

In this section, we integrate Budget-Based Access Control Model with the FUB access control system. Our model consists of three modules: Request Evaluation Module, Access Control Manager Module, and Price-Payment Module. Figure 1 provides an overview of an access control decision.

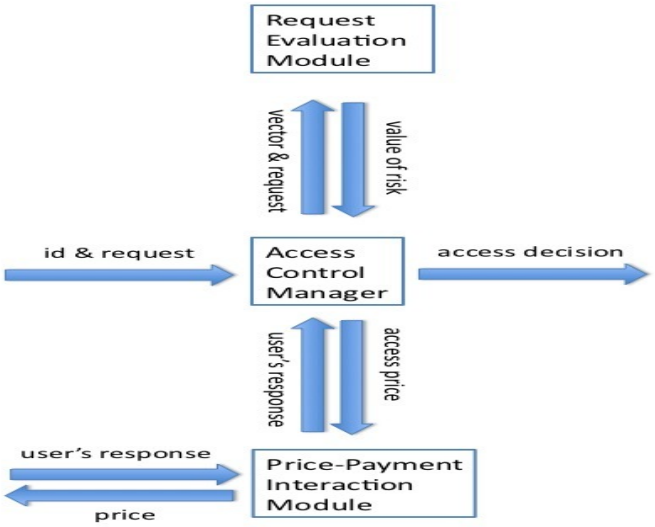

Figure 1: Budget-Based Access Control Model

Each user in this system is assigned a role, which is defined as a combination of his official position and job task. The Access Manager Module runs a batch job to communicate with the human resources system periodically, just as the way the FUB works now, to update each user's role. The users' roles determine their risk budget parameters (risk budget balance, r, lower risk bound, l, and upper risk bound, h).

When a user requests access to an organizational resource or a permission for an activity, he submits a *request* along with his unique *id* to the Access Manager. Access Manager first retrieves that user's information using his *id*. The information is contained in a vector which stores his riskiness, t, and his risk budget parameters (r, l, and h). The riskiness is learned from his behavior history.

The *request* and the retrieved user's information vector are then sent to the Request Evaluation Module, where the risk pricing mechanism works to quantify the risk value of the access request. Risk quantification allows an organization to bound the amount of aggregate risk they are willing to assume.

The access price p is sent back to the Access Control Manager. The Access Control Manager makes access decision according to the subject's risk budget:

- If $p < l$, (Green Zone), access is allowed. User's riskiness t is updated using the learning mechanism.
- If $l < p < h$, (Yellow Zone), an exception needs to be made to grant access to the requester. This can only be done through paying the price of the access. The request is now considered as a request

of exceptional access. Such a price is passed to the Price-Payment Interaction Module, where the risk budget mechanism is launched.

- If $h < p$, (Red Zone), access is not allowed. User's riskiness t is updated using the learning mechanism.

When the access request is within the yellow zone, the Price-Payment Interaction Module is activated with risk budget mechanism. A payment in risk point is required as a proof of the users' willingness to bypass the control with an exception. The user can obtain the exception only if he pays the price. Otherwise, his request is denied.

3.2 *Riskiness Learning Mechanism*

In our model, we employ a reputation model to represent and evaluate each trader's riskiness (t). We integrate two sources for evaluating a riskiness: riskiness based on individual behavioral reputation (t_i) and riskiness based on organizational role-based reputation (t_r). When we integrate our model with a role-based access control system, the initial values of role-based riskiness also comes from the role. For simplicity, we assume the role-based value is manually assigned by system administrator based on subjects' roles, the same as the assignment of risk budget parameters. In the human-subject experimental evaluations, we consider all subjects as having identical roles and assign each of them an role-based riskiness value of 3.

We learn the individual risk behavioral reputation is in a way the evaluation system in eBay works. We consider the initial individual riskiness value is zero:

$$t_{i0} = 0$$

We learn the risk behavior from subjects' access request history. The intuition is to learn users' risk attitude from their behavior: a user always performing low risk activities is considered risk averse; an attempt or request of high risk activity is considered an indicator of suspicious insider who can not be trusted or risk seeking. In principle, every time after a subject submits an access request, the riskiness learning mechanism assigns a discrete feedback value (t_f) to the subject:

- -1 if the request is within green zone;
- 0 if the request is within yellow zone;
- 1 if the request is within red zone.

We used stock trading scenario to evaluate our model in experimental evaluations. We assigned the following feedback value after each trader completed a trade:

- -1 if the daily return is positive for current trading round;
- 0 if he chose not to operate or the daily return is zero;
- 1 if the daily return is negative for current trading round.

The feedback value is then weighted and aggregated to compute the individual riskiness using the following formula:

$$t_{ij} = t_{ij-1} + t_{fj} \times \omega_j$$

ω_j is the weight of the feedback value of j^{th} access request. After the j^{th} trade, the trader's individual riskiness t_{ij} is updated with the previous individual riskiness t_{ij-1} and a weighted feedback value t_{fj}. The latest feedback usually weights more than older ones in practice. For simplicity in this paper, we assumed the feedback values have the same weight ($\omega_j = 0.2$ in our experiments).

A more risk averse user is represented by a smaller value of riskiness. Starting from the initial role-based riskiness, a subject's risk posture is updated with the aggregated feedback value of previous request history:

$$t_j = t_r + t_{ij}$$

Similar to the FUB system, we store the users' riskiness values in human resource database. Our access control model runs a batch job to communicate with the human resource database and update the user's information data periodically. This learning process takes place in the Access Control Manager Module.

3.3 *Access Pricing Mechanism*

Every access right in an organization is priced in risk points. Risk quantification allows an organization to bound the amount of aggregate risk it is willing to assume. In our model, the price of each access (p) provides two functions. First, the price reflects the organizational risk. Second, the price reflects users' willingness to obtain the access. Our model can then leverage the price as a tool to align users' risk behavior with that of the organization.

We first classify the risk level of every access. While information security does not yet have mature metrics for calculating risks, other industries do have well established risk metrics. Many organization are able to assign sensitivity levels to access rights based on a rough estimate of their value. Typically, sensitivity levels correspond to order of magnitude of risk.

Especially, financial institutions are long established uses of risk calculating and management for particular transactions. For example, the Bank for International Settlements (BIS) has sought to standardize regulations and risk calculations for banks internationally [19].

With the estimated risk classification, we use the following formula to calculate the price of an access right when a subject submits a request:

$$p = r_c \times t$$

r_c is the risk classification of that access right, and t is the requester's riskiness. Intuitively, when two users are asking for a same exceptional access, we would like to ask the user with a low riskiness for a lower price while the user with a high riskiness for a higher price. The pricing process takes place in the Request Evaluation Module.

3.4 *Risk Budget Mechanism*

When a requested access is considered as an exception, the risk budget mechanism is launched in the Price-Payment Module. The principles[1] of the risk budget mechanism are as follows:

- Every user is assigned a budget of risk points, a lower risk bound and a upper risk bound based on their roles determined by the organization.

[1]The risk budget mechanism is an instantiation of our theoretical model from [9].

- Users are refused the accesses with prices above their upper risk bound, and granted the accesses with prices below their lower risk bound. When an access's price is between their lower and upper risk bounds, that access is considered as an exception to bypass controls. Users have to pay for the exceptions they request.

- A user is punished if his risk budget is exhausted before his work is completed.

- The more points remain the greater the user's rewards when his work is completed.

- Users' risk budgets are periodically reset.

The requirement of consuming risk points, together with the punishments and the rewards, shifts the cost of risk to the users. The risk budget mechanism illustrates the cost to the users and produces incentives that motivate users to avoid risky activities thus avoid access abuse.

3.4.1 *Risk Budget Assignment*

The risk budget parameters (risk budget balance, r, lower risk bound, l, and upper risk bound, h) are determined by the organization based on the user's roles. In our model, users' roles are manually assigned by the organization, and reflect their official position and job tasks. For instance, an employee who visits rating sites and social network sites to manage the company's reputation will have a large risk budget. An employee in human resources who can access the payroll database will have a very small risk budget. A user with high security clearance may have higher risk bounds than a user with low security clearance.

3.4.2 *Points Payment*

Payment requirements in risk points are generated in the Request Evaluation Module. When the access request within the subject's risk bounds as defined by the yellow zone, payment is required as a proof of the user's willingness to obtain the exceptional access and bypass the controls. The user can obtain the access only if he pays the price; otherwise, his request is denied:

- If the user decides to pay, the access request is allowed, and the user's risk budget balance r is updated.

- If the user refuses to pay, the access request is denied, and the risk budget balance r remains unchanged.

3.4.3 *Punishments*

The incentive in the form of the punishments inflicted on the users refers to some forms of cost that is enforced by the organization and triggered by risk budget exhaustion. It could be an audit or mandatory training program or a loss of access. The budget size implies a risk limit that the organization could bear for a specific individual. And the punishment translates the exhausted budget into a cost that directly aligns the company's and user's incentives. The risk budget connects the risk suffered by the organization and the posted cost borne by the users. As a result, the risk points spent by a user can reflect his willingness to launch a risky action. In the human-subject experimental evaluations, we punish the participants with a loss of access, which prevents them from trading.

3.4.4 *Rewards*

The punishment caused by an exhausted risk budget provides an incentive for the users to avoid risky behavior. However, such a punishment only happens when the users exhaust their risk budgets, which can be rare. It is desirable that the user be encouraged to choose the least risky path for accomplishing his task at all times not only on the edge of budget exhaustion. This strategy minimizes the organization's risk exposure. To this end, we take a measure that rewards the user according to the surplus of his risk budget. Simply speaking, the fewer risk points consumed the more rewards the user will get. In practice, the rewards can be paid in the form of welfare. For example, the user can redeem his unspent points in exchange of a vacation, a bonus, or a prize. In our experiments, we reward the participants a bonus in addition to their compensation of participation. Personalization of pricing also reinforces this feedback.

4 Experimental Evaluation

We conducted two human-subject experiments in order to evaluate our Budget-Based Access Control Model. The first experiment was designed to understand users' behaviors. The second experiment aimed at studying the change of these behaviors within our budget-based access control system, and examining the efficacy of managing operational risks caused by insiders. Section 5 elaborates the outcomes of these experiments.

As discussed in previous sections, the long established use of risk calculation makes financial institutions attractive for deployment of our access control model. Therefore, both experiments mimic stock trading activities. The only difference is that the trading decisions made by the participants in the first experiment are controlled only by their given trading power, while in the second experiment the participants are also restricted by their risk budgets.

It's a common methodology in finance research to use laboratory experiments to study trade-related activities. For example, [20] uses students as experimental subjects. They set up a two-dealer market, in which stock is traded for five rounds and the stock price is labelled with a small number of discrete values. [21] examines the stock market in an experimental multiple-dealer market in which seven actual dealers trade a single stock. We used eight weeks of real market information and stock prices to mimic a single-stock market, in which the participants are asked to maximize their investment return by trading a stock. The details of the experimental design are given in this section.

4.1 *Recruitment*

We recruited ten participants for the experiments and divided them randomly into two groups: five for the first experiment and the other five for the second experiment. All participants were recruited from the Kelley Business School at Indiana University, Bloomington. All of them have stock trading experience. We plan to recruit 90 more participants in our future study.[2]

Although the number of subjects is relatively small in current phase of our study, previous research in finance shows that the experiments with a small number of human subjects are effective in studying human trading behaviors. For example, [20] used two participants to play the active traders and the other two participants to play the market makers, while [21] used seven people to participate in a stock-trading experiment.

[2]The larger study should be completed by the time of the conference.

4.2 *Task Descriptions*

We used eight weeks of historical data to mimic a single-stock market. The first five weeks of the data are used to illustrate the investment environment including the market information and the stock price. The market information are the Nasdaq Index, the S&P 500 Index, the Dow Jones Index. We used the next three weeks of data to simulate 15 trading rounds. Each participant was asked to trade one stock for these 15 rounds, using data for 15 actual trading days with only one round per day. Every round, the participants made their trading decision based on the updated historical figures and a table containing the information from the previous trading day.

The participants practice for the first five rounds to get familiar with the experimental interface. The formal experiment starts on the first day of the seventh week and lasts for 10 rounds. Figure2(a) and Figure 2(b) illustrate the information given to each participant in the first experiment and the second experiment, respectively.

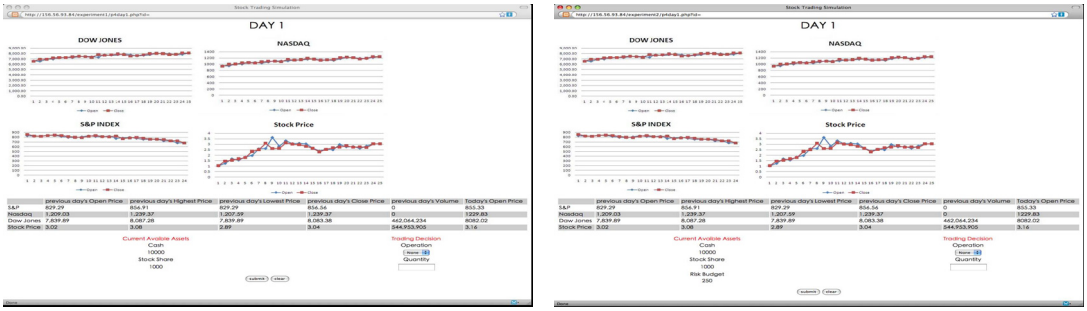

(a) Experiment One (b) Experiment Two

Figure 2: Human-Subject Experimental Evaluations

4.3 *Experiment One*

In the first experiment, the participants begin with an initial trading power of $10,000 cash plus 1,000 stock shares. Given the five weeks of historical data and the information from the previous day, the participants make decisions to maximize their return by trading stock. The date from the previous day contains the following information: open price, highest price, lowest price, closing price, and volume. Also shown is today's opening price, for both the market as a whole and the stock price. The participants can buy, sell or do nothing. Trades will be executed at the opening price of that day. The daily return is calculated using the number of shares traded to multiply the difference between the closing price and the opening price. Each participant's compensation for participating in the experiment is based on that individual's gross return. For simplicity, a participant will get a $10 gift card if he makes a profit at all, or he will get a $5 gift card if he has negative return. Thus participants have the incentive to maximize their return in the experiment.

The participants can only trade within the trading power defined in the experiment. For example, when they spend all their cash they can buy no more stock; when they don't hold any stock share they can not sell stock. This control mechanism simulated the risk management mechanism used in the French bank Societe Generale, whose risk managers only focused on net trading limits. In Experiment Two, we simulate a trading environment under our Budget-Based Access Control.

4.4 *Experiment Two*

In the second experiment, for simplicity, every participant is considered as having the same official position and job task. Therefore, each of them has the same role and an initial risk budget with identical parameters (risk budget balance, r=500, lower risk bound, l=10, and upper risk bound, h=300). Participants begin with 500 risk points and a trading power of $10,000 cash and 1,000 shares of stock. The trading power is the same as in Experiment One. Given the same historical information, they make decisions to maximize their return. The daily return is similarly measured using the number of shares traded to multiply the difference between the closing price and the opening price. Unlike the participants in Experiment One, the participants in the second experiment are compensated based on their gross return and their risk budget balance. For example, if a participant makes a profit or a loss, he gets a $10 or a $5 dollars gift card plus a bonus, respectively. The bonus is calculated such that a participant receives $5 if his risk budget balance is greater than 250, while he gets no bonus if his risk budget balance is less than 250.

As we described in previous sections, the risk budget mechanism also embeds a punishment. In the second experiment, the punishment is provided in the form of a loss of access. For example, a participant is not able to buy or sell any stock if his risk budget is exhausted. In this way, the control mechanism in Experiment Two produces incentives that motivate the participants to simultaneously maximize their return and control their risk.

4.5 *Risk Classification*

In finance, there are various complicated formulas to estimate and quantify the operational risk in stock trading. For simplicity, we adopt the following method to classify the participant's trading requests according to [22]. The principles are as follows:

- a trade is low risk if its potential loss is less than 2% of the trading capital;
- a trade is medium risk if its potential loss is between 2% to 5% of the trading capital; and
- a trade is high risk if its potential loss is greater than 5% of the trading capital.

In this paper, the potential loss caused by a trade is calculated by the number of shares traded times the largest stock price change in previous trading days. We then evaluate each trading request in risk points with different magnitudes (r_c):

- a low risk trade costs 1 risk point;
- a medium risk trade costs 10 risk points;
- a high risk costs 100 risk points;

4.6 *Riskiness*

Compared to other risk-based access control systems [1][2][3][4], a significant difference our Budget-Based Access Control Model has is we learn users' riskiness from their past behaviors, and use the riskiness to update the risk price continuously.

In both experiments, we use a scale of 5 to measure the riskiness (t), where a value of 1 represents the most risk averse user while a value of 5 represents the most risk seeking user. We assume every participant has a role-based riskiness value (t_r) of 3, a zero initial individual riskiness ($t_{i0} = 0$), and that every feedback value (t_f) has an equal weight ($\omega = 0.2$) in the computation of the individual riskiness

(t_i). Each participant then begins his experiment with a t value of 3 and is updated every round using a simple rule:

- the value of t is decreased by 0.2 if the daily return is non-negative for current trading round; and
- the value of t is increased by 0.2 if the daily return is negative for current trading round.

4.7 *Price of Trades*

The price to execute a trade is determined using the value of t to multiply the risk magnitude r_c. For example, a participant with a riskiness value of 3.4 submits a trade request. Based on his trading capital, that request is classified as a medium risk trade. Thus, the cost to complete this trade is 34 risk points. According to his risk budget parameters of that time (e.g., r=380, l=10, h=300), this request is within the yellow zone and a risk communication will be launched for a payment confirmation.

We expect 3 variables to evaluate our Budget-Based Access Control Model: the difference in gross returns, a difference in the time used to complete each experiment, and a difference in the risk exposure. The experimental results are concluded in the next section.

5 Data Analysis

We recorded the experimental results. The results of Experiment One consists of each participant's trading decisions, the gross return each participant made and the time each participant required for a trading decision. The participants make their trading decisions based on the given historical information. Their goal are to maximize their gross return. The trading decisions made by the participants in Experiment One are recorded in Table 1 in Appendix.

The participants in Experiment Two made their trading decisions based on the same given historical information. Unlike in the first experiment, the participants' goal was not only to maximize their gross return but also to minimize the risk caused by their trades. The results of Experiment Two contain the same information of the first experiment plus the risk budget balance for each participant after each round. The trading decisions made by each participant in Experiment Two are recorded in Table 2 in Appendix.

5.1 *Risk Exposure*

We first consider the risk exposure in each experiment. Each participant begins his experiment with an initial risk budget of 500 risk points. The consumption of the risk budget is considered an indicator of the organizational risk exposure. Using the recorded trading decisions and the same risk calculation, we first compute the risk points used for each participant in every round in Experiment One. The risk budget balance that would have been calculated for each participant in Experiment One are recorded in Table 3. Table 4 contains the risk budget balance for each participant as calculated in Experiment Two.

The average consumption of risk points to complete the 15 trading rounds is 108 in Experiment One, while this number is 89 in Experiment Two. Our Budget-Based Access Control Model is designed to regulate aggressive and risky activities, control organizational risk, prevent privileged access abuse, and hence mitigate insider threats. The human-subject experiments were conducted to illustrate how our model could influence the aggressive traders.However, Figure 3 compares the average balance of the remaining risk budgets when the experiments were completed. It shows that our Budget-Based Access Control Model did reduce the organizational risk caused by the trading activities, even when the users are not very aggressive. This result also verifies the conclusion in [9] that the risk budget approach can

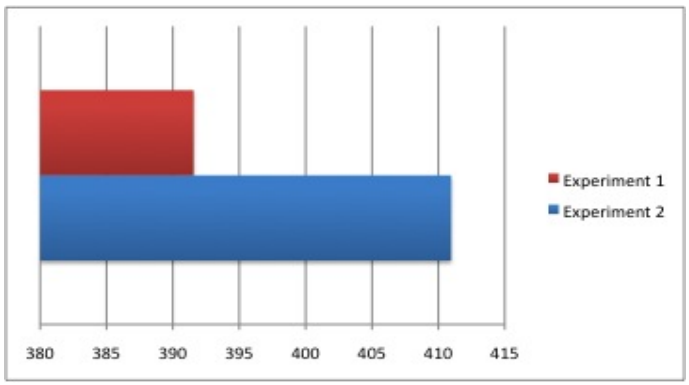

Figure 3: Remaining Risk Budget Balance in Two Experiments

shift the users' risk behaviors from a strong preference for risk seeking to a strong preference for risk aversion.

5.2 *Adoption Cost*

We now focus on the adoption cost of our Budget-Based Access Control. In order to complete the experiment, the participants need 192 seconds on average in the second experiment, while the participants in the first experiment need 138 seconds. In other words, there is a cost of 39% of required time to adopt our model. Figure 4(a) shows the time used to complete each experiment.

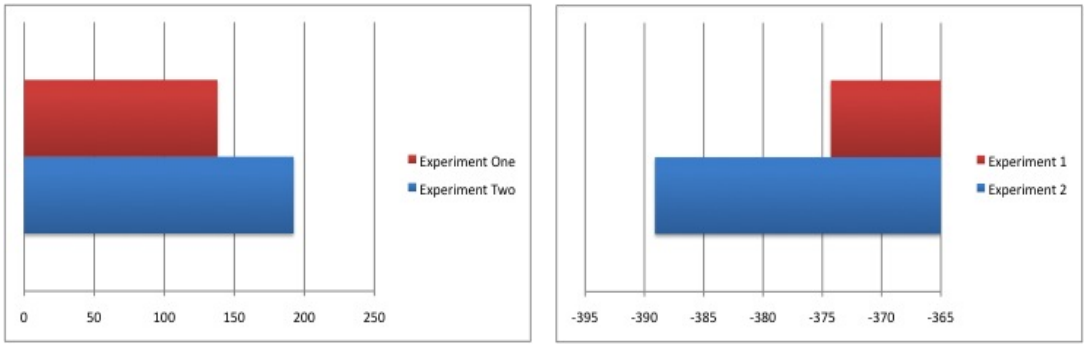

(a) Time Used (in seconds) to Complete Each Experiment (b) Gross Returns in Each Experiment

Figure 4: Adoption Cost

Another measurement of adoption cost is computed using the gross return. The average gross return in Experiment One is a loss of $374 and the number is a loss of $389 in Experiment Two. We noted there is market risks throughout the experimental evaluations. But both groups bore the same market risks. We believe the difference of the gross return can be used to measure the adoption cost. Thus the adoption cost is then measured as $15 or 4% of gross return. Figure 4(b) compares the gross return in each experiment.

We consider the adoption cost as the indicator to parameterize our access control model. The incentives consisting of punishment and reward must compensate the adoption cost in order to make our Budget-Based Access Control Model feasible and favorable. The organization can choose its own will-

ingness to pay in order to avoid risk.

The experimental results show that none of the participant in Experiment Two ignored his risk budget. They are willing to be rewarded by managing their trading risk, and hence to take the adoption cost. It shows that a cost of a 39% of time or a 4% of gross return could be easily compensated by the punishment and reward we introduced in our experiments. This suggests that our model does incur adoption cost but that these are manageable.

6 Conclusion and Future work

Insiders pose a grave threat to the security of organizations. We proposed a Budget-Based Access Control Model to allow access exceptions, in the mean time to mitigate the insider threat caused by these exceptional access. Our model assigns individual users a risk budget that represents the amount of risks an organization can tolerate from that employee. Each action of a user will cost him certain risk points. If the budget is depleted and the user did not finish his work, a penalty ensues. On the other hand, those who diligently reduce the organization's risk, as manifested by the surplus of their budget, will be rewarded. Our experimental study shows that our approach exerts positive influence on rational users' risk attitudes, and reduces the organizational risk caused by the access exceptions. Our approach offers organizations the ability to better fine tune their willingness to pay for risk mitigation. In the future, we plan to study the effectiveness of our approach. We expect to improve our models not only as a technology but also an investment.

References

[1] A. Yemini, D. Dailianas, Florissi, and G. Huberman. Marketnet: Market-based protection of information systems. In *The 12th Int. Symp. on Dynamic Games and Applications*, 2006.

[2] MITRE Corporation. Horizontal integration: Broader access models for realizing information dominace. Technical Report JSR-04-132, JASON Defense Advisory Panel Reports, 2004.

[3] Pau-Chen Cheng, Pankaj Rohatgi, Claudia Keser, Paul A. Karger, Grant M. Wagner, and Angela Schuett Reninger. Fuzzy mls: An experiment on quantified risk-adaptive access control. In *IEEE Symposium on Security and Privacy*, pages 222–230, 2007.

[4] I. Molloy, P. Cheng, and P. Rohatgi. Trading in risk: Using markets to improve access control. In *New Security Paradigms Workshop*, Olympic, California, September 2008. Applied Computer Security Associates.

[5] Insider may have breached more than 10000 patient records at johns hopkins, May 2009.

[6] G. Wearden. The biggest rogue traders in history, January 2008.

[7] Joel Predd, Shari Lawrence Pfleeger, Jeffrey Hunker, and Carla Bulford. Insiders behaving badly. *IEEE Security and Privacy*, 6(4):66–70, 2008.

[8] Nalneesh Gaur, Jason Gaswirth, and Linda Najim. Notes on a scandal: Lessons in operational risk management from societe generale. Technical report, Diamond Management and Technology Consultants, Inc., 2008.

[9] Debin Liu, XiaoFeng Wang, and L. Jean Camp. Mitigating inadvertent insider threats with incentives. In *The proceeding of Financial Cryptography and Data Security*, February 2009.

[10] Mario Scalora and Denise Bulling, editors. *Developing Threat Assessment Best Practice Standards: Leveraging Behavioral Science Strategies to Enhance Decision-Making*, February 2007. Open source literature review in partial fulfillment of Subcontract Number: MSMA-07-00001.

[11] M. Keeney, E. Kowalski, D. Cappelli, A. Moore, T. Shimeall, , and S Rogers. *Insider threat study: Computer system sabotage in critical infrastructure sectors*. May 2005.

[12] Nam T. Nguyen, Peter L. Reiher, and Geoffrey H. Kuenning. Detecting insider threats by monitoring system call activity. pages 45–52. IEEE, 2003.

[13] Ramkumar Chinchani, Anusha Iyer, Hung Ngo Q., and Shambhu Upadhyaya. A target-centric formal model for insider threat and more. October 2004.

[14] Brad Wood. An insider threat model for adversary simulation. In R. H. Anderson, T. Bozek, T. Longstaff, W. Meitzler, M. Skroch, and K. V. Wyk, editors, *Research on mitigating the insider threat to information systems*, Pittsburg, PA, 2000. RAND Corporation.

[15] D. M. Cappelli, A. G. Desai, A. P. Moore, T. J. Shimeall, E. A. Weaver, and B. J. Willke. Management and education of the risk of insider threat (merit): System dynamics modeling of computer system sabotage. In *The International Conference of the System Dynamics Society*, Nijmegen, The Netherlands, July 2006.

[16] Marisa Reddy Randazzo, Dawn M. Cappelli, Michelle M. Keeney, Andrew P. Moore, and Eileen F. Kowalski. Insider threat study: Illicit cyber activity in the banking and finance sector. Technical report, 2004.

[17] Farzaneh Asgharpour, Debin Liu, and L. Jean Camp. Mental models of security risks. In Sven Dietrich and Rachna Dhamija, editors, *Financial Cryptography*, volume 4886 of *Lecture Notes in Computer Science*, pages 367–377. Springer, 2007.

[18] Andreas Schaad, Jonathan Moffett, and Jeremy Jacob. The role-based access control system of a european bank: A case study and discussion. In *In Proceedings of the Sixth ACM Symposium on Access Control Models and Technologies*, pages 3–9. ACM Publish, 2001.

[19] Basel Committee on Banking Supervision. International convergence of capital measurement and capital standards. Technical report, Bank for International Settlements, Basel II, June 2006.

[20] Robert J. Bloomfield and Maureen O'Hara. Market transparency: Who wins and who loses? *Review of Financial Studies*, 12(1), 1998.

[21] Mark D. Flood, Ronald Huisman, Kees C.G. Koedijk, and Ronald J. Mahieu. Quote disclosure and price discovery in multiple dealer financial markets. *Review of Financial Studies*, 12(1), 1998.

[22] Day trading risk management, 2009.

Debin Liu is a PhD candidate at School of Informatics and Computing, Indiana University, Bloomington, Indiana since Fall 2005 with major in Security Informatics and minor in Finance and HCI. His study focuses on Usable Security and Risk Management. His PhD advisor is Dr. L. Jean Camp. Debin received his M.S. degree in Physics from Texas AM University, College Station, Texas in 2005. Before the study in United States, Debin received his B.S. degree in Modern Physics from University of Sci. Tech. of China, Hefei, China in 2003.

Jean Camp is an Associate Professor at the School of Informatics and Computing, Adjunct Professor of Telecommunications, and an Adjunct Professor of Computer Science at Indiana University. Professor Camp's core interest is technical trust mechanisms in economic and social context. It was this interest that led Prof. Camp from graduate electrical engineering research in North Carolina to the Department of Engineering and Public Policy at Carnegie Mellon, and it remained her core interests as a Senior Member of the Technical Staff at Sandia National Laboratories. At Sandia National Laboratories her work focused on computer security. She left Sandia National Laboratories for eight years at Harvard's Kennedy School. Now as a tenured Professor at Indiana University's School of Informatics her research addresses security in society.

XiaoFeng Wang received his Ph.D. in Computer Engineering from Carnegie Mellon University in 2004. He joined Indiana University at Bloomington as an assistant professor in 2004. His research interests span all areas of computer and communication security. Particularly, he is carrying out active research on system and network security (including malware detection and containment, countermeasures to denial of service attacks), privacy-preserving techniques and their application to critical information systems (such as medical information systems), and incentive engineering in information security.

Lusha Wang received her master degree from School of Informatics and Computing, Indiana University in 2010.

Appendix

	Day_31	Day_32	Day_33	Day_34	Day_35	Day_36	Day_37	Day_38	Day_39	Day_40
P_1	sell 100	buy 200	buy 200	buy 200	buy 50	sell 200	buy 300	buy 300	sell 500	buy 300
P_2	buy 500	buy 500	sell 500	sell 500	buy 300	buy 300	buy 50	sell 350	sell 300	buy 100
P_3	buy 30	buy 10	no trade	sell 15	sell 5	no trade	buy 20	sell 40	sell 20	sell 5
P_4	buy 300	sell 200	sell 300	no trade	buy 500	buy 300	sell 200	sell 400	sell 150	buy 400
P_5	sell 1000	no trade	buy 100	buy 100	buy 100	no trade	no trade	buy 100	no trade	buy 400

Table 1: Trading Decisions in Experiment One

	Day_31	Day_32	Day_33	Day_34	Day_35	Day_36	Day_37	Day_38	Day_39	Day_40
P_6	buy 100	sell 222	buy 300	buy 100	sell 300	buy 100	buy 250	buy 140	sell 200	sell 100
P_7	buy 500	buy 100	sell 300	no trade	buy 200	no trade	buy 100	sell 600	sell 400	no trade
P_8	buy 500	buy 250	no trade	sell 300	no trade	buy 200	no trade	sell 300	sell 300	buy 100
P_9	buy 500	sell 200	no trade	no trade	sell 150	buy 150	buy 200	sell 500	no trade	buy 100
P_10	buy 200	buy 200	no trade	no trade	no trade	buy 200	buy 200	no trade	sell 1800	no trade

Table 2: Trading Decisions in Experiment Two

	Day_31	Day_32	Day_33	Day_34	Day_35	Day_36	Day_37	Day_38	Day_39	Day_40
P_1	496.8	493.4	489.8	486.4	483.2	480.2	477.4	474.4	471.2	468.2
P_2	497.2	467.2	464.4	434.4	431.6	428.6	425.8	423.2	420.4	417.8
P_3	497	493.8	493.8	490.4	486.8	486.8	483.4	480.2	476.8	473.2
P_4	497.2	494.6	470.6	470.6	468.4	466	463.4	439.4	436.8	434.4
P_5	180	180	176.6	173.4	170.4	170.4	170.4	167.2	167.2	164.2

Table 3: Risk Balance in Experiment One

	Day_31	Day_32	Day_33	Day_34	Day_35	Day_36	Day_37	Day_38	Day_39	Day_40
P_6	497.2	494.6	491.8	489.2	486.4	483.4	480.6	477.6	474.4	471
P_7	497.2	494.2	491.4	491.4	488.8	488.8	486.4	464.4	440.4	440.4
P_8	497.2	494.2	494.2	491	491	487.6	487.6	484.4	481	477.8
P_9	497.2	494.6	494.6	494.6	491.8	488.8	486	460	460	457.6
P_10	497.2	494.2	494.2	494.2	494.2	491	488	488	208	208

Table 4: Risk Balance in Experiment Two