

A Calculus for the Qualitative Risk Assessment of Policy Override Authorization

Steffen Bartsch
TZI – Universität Bremen
P.O.B. 33 04 40
28334 Bremen, Germany
sbartsch@tzi.de

ABSTRACT

Policy override is gaining traction in the research community to improve the efficiency and usability of authorization mechanisms. These mechanisms turn the conventional privileges into a soft boundary that may be overridden by users in exceptional situations. The challenge for the practical deployment of the policy override mechanisms often is whether policy override is adequate and, if so, to which extent. In this paper, we propose a calculus to support this decision-making process. The calculus is based on proven risk assessment practices and derives a qualitative result on the adequacy for specific roles and override extents. Moreover, we developed a tool to support the policy override risk assessment. The calculus and the tool are briefly evaluated in two distinct contexts.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: Security and Protection

General Terms

Security, Management

Keywords

Authorization policy, policy override, risk assessment

1. INTRODUCTION

In *The Use of Knowledge in Society*, Friedrich Hayek argues that in economies, the information for making choices is present on the spot and not centrally where the plans are usually made [16]. Similarly, many decisions about authorization rules are made far from the employees' workplaces where those decisions threaten efficiency. Even if the local employees are involved in the building of authorization rules, the employees typically only take the most common processes into account. In practice, every once in a while, employees need higher privileges than assigned by default to complete the work at hand. In these cases, they either

turn to coworkers with the appropriate privileges, request at the IT department to have the privileges extended or are handed the credentials of someone with higher privileges. All of these approaches have drawbacks, such as inefficiencies or the uncontrollable consequences in the last case.

To overcome these inefficiencies while keeping the organization's risk in check, we need to consider both the *least-privilege* principle [32] as known from conventional authorization schemes as well as a *most-tolerable-privilege* approach with safeguards in place. Researchers have proposed *Optimistic Security* mechanisms, called at times Policy Override, Authorization Escalation or Break-glass mechanisms, to allow a soft boundary at the normal least-privilege extent of privilege and a hard boundary at the most-tolerable-privilege extent [4, 28, 36, 30, 10].

In health care, policy override mechanisms are already in use [12, 21, 23]. Still, it remains difficult for organizations from other sectors to decide which risks are adequate for the estimated gains. This decision is rather straight-forward for hospital health care information systems with the high potential gain of saving life. For other environments, the decision is not as clear-cut because organizations need to balance a possibly increased insider threat against potential gains. When the authorization override model allows fine-grained settings of who may override to which extent, a considerable amount of work needs to be done to analyze risks and gains for each case.

To ease the effort in this decision-making process, we propose a novel calculus in this paper to support the security and domain experts' decision. The Policy Override Calculus estimates the override adequacy by employees' roles and override extent. At the heart of the calculus, the calculus formulae and operators describe how to derive the adequacy from the collected data. The policy override calculus is based on proven risk management practices. The main contribution of the proposed approach is the adaption of the available risk assessment methods for the specific requirements of policy override configuration. The calculus has been implemented in a override risk management tool. The calculus and the tool are evaluated briefly in two distinct contexts: Firstly, by applying a real-world example taken from a medium-sized company, which already employs policy override in its business web application. Additionally, we evaluated the appropriateness of the approach for a larger company.

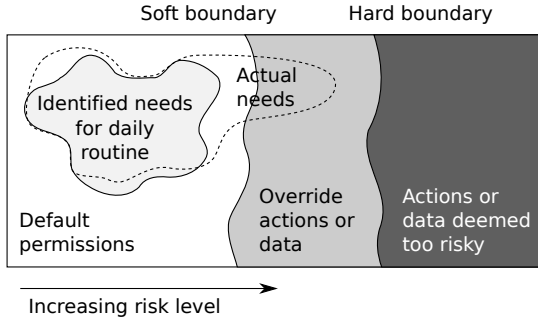


Figure 1: Abstract example of a role's hard and soft boundary

2. POLICY OVERRIDE

When creating the policy for an application's authorization, security experts conduct interviews and process analysis to identify each employee function's needs. Figure 1 shows an abstract example of a role's requirements lightly shaded on the left side. Depending on the effort and analysis process employed, the resulting authorization requirements may or may not be accurate. In the figure, the actual needs are shown with a dotted line.

The next step for security experts is to derive the role's permissions from the authorization requirements. Most likely, the permissions will not exactly match the requirements as authorization models are typically coarse when compared to the identified requirements. The default permissions are shown as white area in Figure 2. In conventional authorization models, the employee has to make do with these permissions. As stated in the introduction, for any missing permissions in daily routine or exceptional situations the employee needs to turn to co-workers with higher privileges. Alternatively, the employee may formally apply for the assignment of the missing permissions. Oftentimes, these exceptional assignments then remain in effect even if caused by a single incident. While these are valid options for some cases, the bureaucracy and inefficiency make these options inadequate for other cases.

When using an application that supports policy override, the employee may have the additional option of overriding the denied permission. As proposed by Cheng et al. [10], we introduce a *soft boundary* as an addition to the hard boundary in conventional authorization models. As shown in Figure 1, the soft boundary separates the default permissions from those that can be gained through overriding the policy. Any action in override will be logged and audited. The hard boundary marks the limit of privileges gainable through policy override. Anything beyond the hard boundary is thought to be of too high risk to allow the specific role to have access to it. Note that it may make sense to disallow any override for certain roles so that the hard and soft boundary run along the same line.

3. POLICY OVERRIDE AUTHORIZATION MODEL

To implement policy override, the authorization model needs to be adapted. We base the authorization model on a con-

ventional RBAC model with hierarchies, $RBAC_1$ as proposed by Sandhu et al. [33]:

- *User, Role and Permission*
- Permission assignments: $PA \subseteq Permission \times Role$
- User assignments: $UA \subseteq User \times Role$
- Role hierarchy as a partial order $RH \subseteq Role \times Role$, with role dominance, denoted by the symbol \geq
- $JuniorRoles : (Role, Session) \rightarrow \mathcal{P}(Role)$, with $(r, s) \mapsto \{r' | r \geq r'\}$
- $SessionUser : Session \rightarrow User$, returning the session's user
- $SessionRoles : Session \rightarrow \mathcal{P}(Role)$, the session's activated roles, with

$$SessionRoles(s) \subseteq \{r | \exists r' [r \in JuniorRoles(r', s) \wedge (SessionUser(s), r') \in UA]\}$$
- $SessionPermissions : Session \rightarrow \mathcal{P}(Permission)$, with

$$s \mapsto \{p | (p, r) \in PA \wedge r \in \cup_{r' \in SessionRoles(s)} JuniorRoles(r', s)\}$$

To implement policy override, another role relation is introduced to define which roles a user may extend his privileges to. This *Override Roles* relation is then used to modify the available roles that can be activated when the session is in override mode:

- *Override roles*: $OR \subseteq Role \times Role$ is a relation that defines to which role the holder of the first role may extend her privilege in case of override
- $IsOverrideMode : Session \rightarrow bool$ is a predicate that indicates whether a user has activated the override mode on a specific session
- $JuniorRoles$ is redefined to return roles in the override hierarchy if override mode is active for the current session:

$$(r, s) \mapsto \begin{cases} \{r' | r \geq r' \vee \\ \exists r'' [(r'', r') \in OR \wedge r \geq r'']\}, \\ \text{if } IsOverrideMode(s) \\ \{r' | r \geq r'\}, \text{ else} \end{cases}$$

The policy override modifications do not affect authorization constraints, which are enforced as usual and are for brevity not included in the definition.

4. POLICY OVERRIDE CALCULUS

When defining an authorization policy with override, security and domain experts not only need to assign the normal privileges of each role. For the normal privileges, experts usually evaluate the authorization requirements from a *need-to-know* perspective. In policies that consider override, a second, *hard* boundary needs to be defined in addition to

the *soft* boundary, which is derived from the need-to-know evaluation. The hard boundary balances the risks of additional, not routinely needed knowledge and abilities of each role against potential advantages of policy override. To help with this decision, the Policy Override Calculus gives an estimate of how appropriate specific hard boundaries for individual roles are. In contrast to what the term “calculus” might suggest, the Policy Override Calculus is not a mathematical calculus in the quantitative sense, but still provides formulae, based on risk management principles, to derive the qualitative results.

The adequacy of override by role and override extent is derived from qualitative data collected from domain and security experts for several inputs of the calculus. We decided against a quantitative assessment because quantitative estimations suggest a precision that is not realistic [37, 27]. For example, humans are very imprecise when estimating frequencies of occurrences [15]. Also, many aspects, such as employee motivation in case of insider threats, are not quantifiable. As a result, quantitative results need to be interpreted. For the policy override calculus, a qualitative and relative risk that allows one to rank the risks is sufficient as a decision support. The calculus estimates the additional risk for specific extents when compared to default authorization of a specific role. The qualitative values that are collected for the input components are either *Normal*, *High* or *Very High*, abbreviated “N”, “H” and “V” in the tables.

4.1 Risk

In risk management, the implementation of possible mitigations are often decided based on a risk/cost relationship. Similarly, we balance risks against potential benefits for the policy override calculus. The reasoning is that the benefit that is lost when not using override can be seen as the costs tied to the mitigation of the risks that are caused by override actions. First, we will consider the risks that an organization is exposed to when granting override options to users. For the policy override calculus, the individual risks for the security objectives confidentiality, integrity and availability are aggregated into a single risk value as it is common in information security risk assessment, see e.g. [37]. Depending on the application, a confidentiality breach may lead to the disclosure of sensible information. When the integrity is compromised, employees may work with manipulated data or miss data. Effects of non-availability are relevant to time-dependent tasks. To derive meaningful results for the override adequacy per user system role and a specific privilege extent, the *Risk* is a function of role and extent:

$$\begin{aligned} Risk(role, extent) = & \\ & \propto (SpecificRisk(Confidentiality, role, extent), \\ & \quad SpecificRisk(Integrity, role, extent), \\ & \quad SpecificRisk(Availability, role, extent)) \end{aligned}$$

Following [20], the aggregation operator (\propto) is defined as the maximum of the individual risks. The individual risks do not need to be weighted in this aggregation as weights are implicitly present in each individual risk’s value through the protection need.

		Protection Need		
		Normal	High	Very High
Threat	Normal	N	N	N
	High	N	H	H
	Very High	N	H	V

Table 1: Specific Risk: (Protection need) \otimes (Threat likelihood)

Specific Risk. For each security objective, confidentiality, integrity and availability, the *SpecificRisk* is calculated for each user role and override privilege extent. Throughout the formulae, sans-serif fonts are used to signal direct input data.

$$\begin{aligned} SpecificRisk(objective, role, extent) = & \\ & ProtectionNeed(objective, extent) \otimes \\ & ThreatLikelihood(objective, role, extent) \end{aligned}$$

In risk models, risk is typically defined as the expected loss resulting from a threat as calculated from the product of an incident’s potential damage and its likelihood [35]. Following this approach for the Policy Override Calculus, the *SpecificRisk* is derived as the individual risk to the security objective from the *Protection Need* as the expected impact and the *ThreatLikelihood* as the likelihood of the incidents (\otimes). The \otimes operator is defined as a look-up matrix shown in Table 1 that follows the Risk-Level Matrix proposed in NIST 800-30 guidelines, in which the risk levels are determined through multiplication [37].

One operand for the calculation of the *SpecificRisk* is the impact of incidents that is approximated through the *Protection Need* as defined in *BSI IT-Grundschutz*, a German information security baseline standard [7]. In the standard, the protection levels are defined as follows:

- *Normal* The impact of any loss or damage is limited and calculable.
- *High* The impact of any loss or damage may be considerable.
- *Very High* The impact of any loss or damage could be of catastrophic proportions threatening the survival of the organization.

The IT-Grundschutz standard defines the protection need levels for the impacts regarding the violation of laws, regulations or contracts, privacy rights, physical injury, the ability to complete tasks at hand, internal and external effects and financial consequences [7].

The protection need values are differentiated according to the exposure extent. A company with several branches may be able to absorb the losses of a single incident at one branch. On the other hand, if the whole company is affected, the damage might be disastrous. The potential damage thus depends on the extent of data exposed. An example for the different risk levels by privilege extent is shown in Figure 2. In the example, the disclosure of only some contracts’ quality data has a much smaller impact than the disclosure of all of the company’s contracts.




Figure 2: Example of risk varying according to the privileges of a role

Threat Likelihood. The second operand for the *Specific-Risk* is the *ThreatLikelihood*, defined as the likelihood of the occurrence of an incident. The *ThreatLikelihood* is calculated for a given security objective and role for the case that a specific policy override extent has been assigned to that role:

$$\text{ThreatLikelihood}(\text{objective}, \text{role}, \text{extent}) = \text{RoleThreat}(\text{role}) \odot \text{OpportunityThreat}(\text{objective}, \text{extent})$$

Two aspects are taken into consideration to arrive at an estimation of the threat likelihood: the threat induced by the personal characteristics of users in specific roles (*RoleThreat*) and the differences in threat caused by different privileges (*OpportunityThreat*). This separation is based on the criminological knowledge on insider threat. The insider threat models in literature separate the aspects of the capability of users, the motivation and the opportunity in the computer systems (CMO model) [34, 41, 39, 40]. The capability is of minor interest for the risk in the case of policy override because the risks are expected to originate from ordinary activities similar to an insider’s daily routine rather than sophisticated attacks. On the other hand, motivation and opportunity of users are of high importance. To facilitate the collection of input data for opportunity and motivational factors with a reasonable effort, these factors are collected by the two separate input types, *RoleThreat* and *OpportunityThreat*. This separation can also be found with Magklaras and Funell, who suggest a similar insider threat model that depends on reasons and system role [22]. Schultz also uses the privilege of users as attributes of insiders in his CMO model [34].

Role Threat. As described above, the *RoleThreat* captures threat likelihood aspects that differ regarding to the system roles. Intuitively, for personal characteristics, experts need to evaluate how trustworthy a group of employees is in a specific role. Magklaras and Furnell define a taxonomy of insider threat considering the user’s system role, reasons of attack and the attacks’ consequences [22]. We follow their approach on the reasons of attack and differentiate *accidental* from *intentional* threats. On the intentional side, following the CMO model, the role threat is foremost about the motivation that causes insiders to harm the employer. The NIST 800-30 risk management guide lists human threat sources [37]. From practice, Randazzo et al. assembled a detailed study on insider threat incidents in the banking and financial sector and describe different motivations [29]. Further motivations are described in a U.S. CERT report on critical infrastructure sabotage, e.g. disgruntled employees,

	Intentional	Accidental
Threat likelihood factors	Trust in employees; Employee Satisfaction; Seniority; Previous issues with employees of specific roles; Deterrents	Frequency of application usage; Application usability; Work environment; Training level
Threats to confidentiality	Selling of data	Inappropriate data handling
Threats to availability, integrity	Intentional data manipulation	Accidental removal, modification of data

Table 2: Likelihood factors and threats related to the role threat likelihood

previous sanctions and personal predispositions that may be mapped to users in roles [24]. According to Cappelli et al. insiders tend to occupy foremost lower-level, non-technical positions [9], another factor that supports the correlation of system roles and differing threat likelihoods. Similarly, a study by the Association of Certified Fraud Examiners of 1134 fraud and abuse incidents, not limited to computer system activities, show that the perpetrator’s position and department affects incident frequencies [2].

In addition to motivational factors, threats from accidental incidents need to be considered. Accidental aspects can be derived from previous incidents for specific roles, the training level of employees in a role and similar aspects. A selection of aspects and typical incidents relating to the system role that security experts need to take into account for the *RoleThreat* estimation is listed in Table 2. The qualitative threat values for the *RoleThreat* correspond to an estimation of the likelihood of incidents. The input value *Normal* stands for highly unlikely, *High* for unlikely and *Very High* for likely that an incident occurs.

Opportunity Threat. The second dimension of the insider threat likelihood is the opportunity that is caused by the extent of privilege in override cases. The *OpportunityThreat* is based on the need to differentiate by privilege as the change of data exposure and differences of possible actions in the system influences the threat likelihood. In the criminological CMO model, opportunity is one of the three primary categories. The Rational Choice Perspective from the Situational Crime Prevention theory states that crimes are deliberate and purposive even though the decisioning may be imperfect [40]. Thus, the motivation might be higher with increased opportunity that causes higher interests towards the data for espionage and higher impact in case of sabotage. As described in the *RoleThreat* section, financial gain motivates most perpetrators [29]. In an insider threat prevention guide, Cappelli et al. list motivations for theft and modification for financial gain and business advantage, stating that insiders acted mostly for financial gain by stealing data to sell or modify data for their own or friends’ profit [9]. Another insider threat aspect is that “opportunity makes

	Threat	Likelihood factors
Confidentiality	Selling data for industry espionage	Data value to competitor; Risks in selling
Integrity, Availability	Data manipulation	Gain from fraud; Value of sabotage to competitor

Table 3: Likelihood factors and threats related to the opportunity threat likelihood

		Role Threat		
		Normal	High	Very High
Opport.	Normal	N	H	V
	High	N	V	V
	Very High	H	V	V

Table 4: Threat Likelihood: (Role Threat) \odot (Opportunity Threat)

a thief” as described for insider threats in the 1994 Audit Commission report [3].

The OpportunityThreat is estimated independently from personal characteristics, which are already considered for the RoleThreat. Moreover, the opportunity threat likelihood depends only on the motivational factors for harmful actions so that only intentional aspects are taken into account. The varying impacts from accidental incidents caused by different privileges are already taken into account for Protection-Need. A selection of relevant aspects of threats to Confidentiality, Integrity and Availability are shown in Table 3.

Threat Likelihood Aggregation. Automatically aggregating the RoleThreat and OpportunityThreat values is a key challenge. The focus lies on the aggregation of the different dimensions of the threat likelihood for specific roles with the values for the different privileges. The proposal of the look-up operator \odot is shown in Table 4. It is based on the assumption that “Opportunity makes a thief” [3]. The starting point is the threat estimation for the specific role, which is then modified according to opportunity effects. For Very High opportunities, even solid employees might be tempted to commit a malicious insider act so that the aggregated threat is High. For High and Very High RoleThreat, the aggregated threat increases with higher opportunities.

4.2 Benefit

While there are potential risks related to allowing policy override, an enterprise may significantly benefit from the increased flexibility. The benefits are estimated with the following formula for users of a role with a specific override extent configuration:

$$Benefit(role, extent) = Frequency(role) \times BenefitPerOverride(role, extent)$$

The benefit is calculated from the frequency of override incidents and an estimate of the benefit that may be achieved on average from each override incident. Following quantitative

		Override Frequency		
		Normal	High	Very High
Benefit	Normal	N	H	H
	High	H	H	V
	Very High	H	V	V

Table 5: Benefit: (Benefit per Override) \times (Override Frequency)

		Effort per Override		
		Normal	High	Very High
Gain	Normal	N	N	N
	High	H	N	N
	Very High	V	H	N

Table 6: Net Gain per Override: (Gain per Override) $-$ (Effort per Override)

calculations, the \times operator acts similar to multiplication as shown in Table 5.

Frequency. The daily routine of system users often varies significantly according to the role. In the same way, each role may have different frequencies of situations where quick responses are needed. Thus, domain experts estimate the frequency of override cases by role. Aspects to consider should include the structuredness of daily routine, the frequency of unforeseeable incidents and whether there is direct customer contact.

Benefit per Override. The second component for the calculation of the benefit is the benefit that may be gained in each override incident:

$$BenefitPerOverride(role, extent) = EfficiencyGainPerOverride(role, extent) - EffortPerOverride$$

BenefitPerOverride follows the intuition that each case of policy override brings benefit, but also causes auditing effort for the superior who is responsible for auditing the case. Depending on the auditing implementation, auditing effort may reduce the benefit from policy override cases, particularly with high numbers of policy override cases. Accordingly, BenefitPerOverride derives the net gain that is achieved in each override case from the efficiency gain in each case and the auditing effort per case ($-$). Effort estimated as *High* and *Very high* reduces the net gain similar to subtraction as shown in Table 6.

Efficiency Gain per Override. The efficiency gain per override action is estimated separately for each role and permission extent. In the example of a quality operator in Figure 2, domain experts may foresee cases in that the company may profit to a larger extent from the access of data of the whole assigned branch because the operator might jump without any bureaucracy from one local contract to another. On the other hand, the experts may find it unlikely that the same employee would need to switch branches quickly. Aspects

		Aggregated Risk		
		Normal	High	Very High
Gain	Normal	N	L	L
	High	H	N	L
	Very High	V	H	N

Table 7: Override Adequacy: (Aggregated Risks) \bowtie (Net Gains)

to be considered are the company gain per override, the time saved by not requiring a formal delegation of permissions and the likelihood of work in the context of a specific extent. The qualitative input values should only be above *Normal* if there is additional gain when compared to the original privilege extent of a role.

Auditing Effort. The additional effort that the company needs to invest per override is caused by the need to audit override actions afterwards. For the Policy Override Calculus, domain experts estimate the typical effort that is spent on auditing the actions per override.

4.3 Override Adequacy

From the risk and benefit estimations, the policy override adequacy is determined as decision support for policy override authorization configuration for each role with a specific override extent assigned:

$$Adequacy(role, extent) = Risk(role, extent) \bowtie Benefit(role, extent)$$

The \bowtie operator balances risks with benefits. The interpretation of the outcome of this operator depends on the company policy with regard to acceptable risks. The U.S. FIPS 191 guideline suggests to calculate risk/cost relationships for the balancing of security mechanism costs against risks based on qualitative data [26]. Similarly, the look-up table for the operator definition given in Table 7 is derived from the risk/benefit relationship. In this case, the costs are the lost benefits from not employing override as a way to mitigate the risks from override actions. The adequacy values are calculated by quantifying Normal as 1, High as 2 and Very High as 3. The result from the ratio $a = Gain/Risk$ is interpreted as Low for values $a < 1$, Normal for $1 \leq a < 1.5$, High for $1.5 \leq a < 2.5$ and Very High for values $a \geq 2.5$. Results from this operator do not offer an absolute estimation of the override adequacy. Rather, the results help by providing an order of the most suitable role/privilege extent combinations.

5. EARLY EVALUATION

For the evaluation of the proposed calculus, we developed the Web-based Override Risk Assessment tool to facilitate the collection of input data and calculation of the override adequacy, depicted in Figure 3. In the override risk tool, input tables are provided for each of the input types. For each input type, a separate tab is given in the top area, in which experts can enter input data. If necessary, even the operators may be modified here. The resulting override adequacy is continuously calculated and displayed in the bottom area, so that stakeholders who provide the input for the evaluation can directly see the consequences of inputs. Coloring

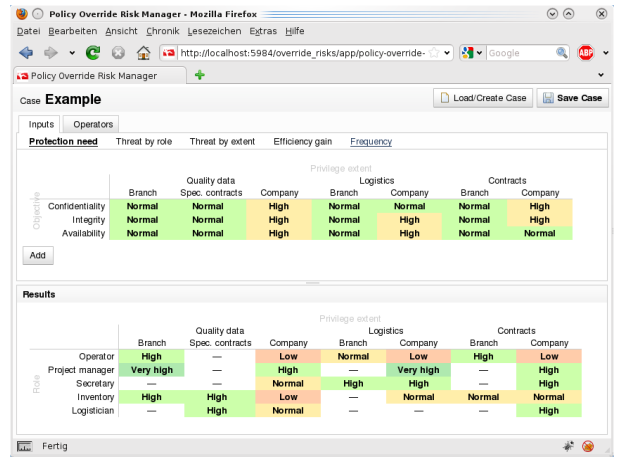


Figure 3: Override Risk Assessment tool

according to the qualitative value eases the collection of data and interpretation of results.

We have evaluated the proposed override calculus in two distinct environments as case studies. The context of the first evaluation is a medium-sized enterprise that has implemented a business application with policy override. The business application supports the enterprise’s quality management processes. With policy override in place, the employees have already built up experience with the policy override concepts. To evaluate the developed calculus, we worked together with the company’s manager to collect the necessary data. Parts of the override adequacy results from the analysis with the Override Risk Assessment tool are shown in Table 8 by role and privilege extent. The privileges correspond to application areas in the authorization policy. The following roles of internal employees from the authorization policy are shown. In the *Inventory* role, the employees manage stock in the warehouse. *Logisticians* are office workers concerned with logistics. *Operators* perform the quality control tasks. *Secretaries* and *Project managers* fulfill administrative and management duties. Because of space considerations, the individual input values are not shown in this paper.

As shown in the result table, the calculus only outputs values for permissions not already assigned for the conventional *need-to-know* permission. The results are very heterogeneous. For example, the results strongly suggest policy override for *Project managers* with two *Very high* and two *High* results. *Operators*, on the other hand, should only receive override privileges for some extents as shown by the *Low* values. After the data collection, we interviewed the manager from the studied company on the plausibility and relevance of the results. The manager agreed with the results and decided to adapt the original override policy. The override calculus thus supported the domain experts’ intuition in the complex evaluation of risks and gains.

The second evaluation context is a large enterprise from the food industry with about 1,000 employees. In this case, an IT management person used the Override Risk Assessment tool with assistance. While the first evaluation focused

Override Adequacy	Privilege Extent						
	Logistics		Contracts		Quality data		
	Branch	Company	Branch	Company	Spec. contracts	Branch	Company
Inventory		N	N	N	N	N	L
Logistician				H		H	N
Operator	N	L	H	L		H	L
Secretary	H	H		H			N
Project manager		V		H		V	H

Table 8: Example result for the override adequacy by role and privilege extent

on usefulness of the adequacy results, the second case was aimed at the usability of the calculus. Primarily, the evaluation looked at the necessary effort and the understandability of the involved risk and benefit concepts. For that reason, the policy override calculus was only collected for part of the company’s roles and an interview was conducted on the studied context and the participant’s impressions. In the studied environment, no explicit risk assessment had been conducted for authorization policy decisions before. Configuration had only been evaluated informally. Thus, the estimation of protection needs and threat likelihoods were initially perceived as a challenge. Still, the concepts that are used in the calculus could be applied. One minor issue was that the available levels to choose from are abstract and natural-language categories instead of “Normal” and “High” could improve the process. In the company’s enterprise system, only a flat role structure has been implemented, resulting in a large number of roles. The participant thus saw a high effort required to fully analyze the company with the policy override calculus. Here, similar to the initial role engineering, external support might be necessary. On the other hand, he was unsure whether the role-based distinction was fine-grained enough to estimate personal threat likelihoods.

From the two small-scale evaluations, we draw the conclusion that the current design of the calculus and implementation of the Override Risk Assessment tool may be helpful in small to medium enterprises. For large enterprises, the effort may be high, at least when there is no structured risk assessment data available. Still, further research and larger-scale evaluations are necessary to show the broader validity of the proposed method.

6. RELATED WORK

To the best of our knowledge, there have been no publications on decision support for policy override authorization configuration. Risk assessment has been included in override authorization models by Cheng et al. in their optimistic authorization model [10]. They use risk quantification for risk/benefit analysis with “risk credit lines.” Without explicitly offering override functionality, the risk-based access control model RAdAC balances operational needs versus security risk [11]. Britton and Brown analyzed risk factors to be used in the RAdAC model [6]. Similarly, Diep et al. described an authorization model with context-based decisions that includes a quantitative risk assessment on each action [13]. For a similar mechanism, Dimmock et al. suggest a Prolog-based risk and trust decision-making mechanism [14]. These models use quantitative methods to automatically judge about the risk of actions and, thus, authorization. In practice, quantitative risk assessment is in many con-

texts missing the necessary reliability for automatic decision-making because of insufficient input data quality. Therefore, we see the decision-support approach, as suggested in this paper, as more practical.

Policy override, on the other hand, is an authorization mechanism that is already in use in health care applications. Accordingly, most literature on practical experience with override and its implementation is from health care environments. Denley and Smith as well as Røstad and Edsberg report on the experience from implementing override in addition to other authorization mechanisms in clinical information systems and analyze the audit logs [12, 31]. Both state the problem of a high number of override cases. There is a U.S. HIPAA (Health Insurance Portability and Accountability Act) document on the usage of policy override in health care [17]. Outside of health care institutions, Stevens and Wulf report on an authorization case study including policy override at a steel mill that analyzes inter-organizational cooperation and competition from a CSCW perspective [36]. In none of these works the balancing of risk is explicitly discussed. Povey introduces policy override as *Optimistic Security* and formulates requirements, but does not mention the balancing of risks either [28]. Zhao and Johnson employ a game-theoretical approach to model incentives in policy override authorization [43, 42]. Their model may help in future developments of the calculus that is proposed in this paper by offering estimates of threats from insiders, the usage frequency and company benefits.

Risk assessment is the second major building block of this work. There is surprisingly little work on threat assessment particularly targeted at insider threats. Cappelli et al. suggest to include insider threat in risk assessment processes [9]. On general risk management, there is a wealth of publications available [5, 8]. First of all, there are national and international standards on risk management. The U.S. FIPS-65 standard, withdrawn in 1995, applied the well-known Annualized Loss Expectancy (ALE) quantitative approach [25]. There are further quantitative approaches, such as Value-at-Risk-based ones [18], and methods that use quantitative calculations on broad inputs [19]. The most widely-used methods are of qualitative nature, though. Very well-known is the NIST Special Publication 800-30 [37] that supersedes the quantitative FIPS-65 standard [25]. Other wide-spread approaches are CRAMM [38] and OCTAVE [1]. These standard methods enact very similar processes: The processes start with a qualitative valuation of assets or impacts of incidents. Then, threats and vulnerabilities are identified and categorized with likelihood estimations, again by way of qualitative values. In a risk assessment step, the inputs

are combined into risk estimations and, in the risk management part, mitigations are chosen.

The risk management processes are very powerful methods for general assessment of risks. To derive the appropriate override privileges, the objective of the calculus presented in this paper, the standard processes are too general and coarse, though. In the calculus, we need to differentiate between different roles and privilege extents in the case of authorization. Moreover, the override calculus leaves out unnecessary threats from non-insiders and vulnerabilities, thus reducing the effort. Still, outputs from the standard risk assessment processes may be very helpful to derive the input values of the proposed override calculus.

7. CONCLUSION

In this paper, we propose a novel calculus that estimates the adequacy of policy override for specific roles and extents. The qualitative result is derived from collected inputs, including the protection needs, threats, and benefits. The evaluations have shown that the calculus may be viable for small to medium enterprises for insights into the adequate extent of policy override. Domain experts of a medium-sized company could improve the policy override configuration from the calculus results and gave positive feedback on the usefulness of the results. In large enterprises, the implementation of an override risk assessment is more challenging, though. Still, with policy override concepts gaining traction in the research community, the calculus helps to overcome a major obstacle that hinders a wider adoption apart from health care environments. The calculus gives relative and qualitative result for the often difficult question whether policy override is adequate in a specific setting.

Future works involve evaluating the policy calculus on a larger scale and more intensively in the health care domain. Then, we would like to evaluate re-using the calculus results for the creation of conventional policies, use different levels of abstraction for input collection and experiment with directly employing the calculus result for override permissions without expert intervention. Lastly, we consider re-using data from Information Security Management Systems and previous risk assessment to reduce the effort of collecting the risk data.

8. ACKNOWLEDGMENTS

Many thanks to Dr. Karsten Sohr and Prof. Carsten Bormann for numerous discussions and helpful comments on the subject of this paper. I would also like to thank the employees of the involved companies for their cooperation and evaluation feedback.

9. REFERENCES

- [1] C. Alberts, A. Dorofee, J. Stevens, and C. Woody. *Introduction to the OCTAVE Approach*. CarnegieMellon, Pittsburgh, PA, USA, August 2003.
- [2] Association of Certified Fraud Examiners (ACFE). Report to the nation on occupational fraud & abuse, 2006.
- [3] Audit Commission. *Opportunity Makes a Thief: an Analysis of Computer Abuse*. Audit Commission Publication, London, UK, 1994.
- [4] L. Badger. Providing a flexible security override for trusted systems. In *CSFW*, pages 115–121, 1990.
- [5] R. Baskerville. Information systems security design methods: implications for information systems development. *ACM Comput. Surv.*, 25(4):375–414, 1993.
- [6] D. W. Britton and I. A. Brown. A security risk measurement for the RAdAC model. Master’s thesis, Naval Postgraduate School Monterey, CA, March 2007.
- [7] Bundesamt für Sicherheit in der Informationstechnik (BSI). BSI-Standard 100-2: IT-Grundschutz-Vorgehensweise. Version 2.0, 2008.
- [8] P. L. Campbell and J. E. Stamp. A classification scheme for risk assessment methods. Technical Report SAND2004-4233, Sandia National Laboratories, 2004.
- [9] D. Cappelli, A. Moore, R. F. Trzeciak, and T. J. Shimeall. Common sense guide to prevention and detection of insider threats 3rd edition – version 3.1. Technical report, CarnegieMellon, January 2009.
- [10] P.-C. Cheng, P. Rohatgi, C. Keser, P. A. Karger, G. M. Wagner, and A. S. Reninger. Fuzzy multi-level security: An experiment on quantified risk-adaptive access control. In *SP ’07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 222–230, Washington, DC, USA, 2007. IEEE Computer Society.
- [11] R. Choudhary. A policy based architecture for NSA RAdAC model. In *Information Assurance Workshop (IAW 05)*, pages 294–301, June 2005.
- [12] I. Denley and S. W. Smith. Privacy in clinical information systems in secondary care. *BMJ*, 318(7194):1328–31, May 1999.
- [13] N. N. Diep, L. X. Hung, Y. Zhung, S. Lee, Y.-K. Lee, and H. Lee. Enforcing access control using risk assessment. In *Universal Multiservice Networks, 2007. ECUMN ’07. Fourth European Conference on*, pages 419–424, feb. 2007.
- [14] N. Dimmock, A. Belokosztolszki, D. Eyers, J. Bacon, and K. Moody. Using trust and risk in role-based access control policies. In *SACMAT ’04: Proceedings of the ninth ACM symposium on Access control models and technologies*, pages 156–162, New York, NY, USA, 2004. ACM.
- [15] T. Gilovich, D. W. Griffin, and D. Kahneman, editors. *Heuristics and biases: the psychology of intuitive judgement*. Cambridge University Press, 2002.
- [16] F. A. Hayek. The use of knowledge in society. *American Economic Review*, 35:519–530, September 1945. Reprinted in F.A. Hayek (ed.), *Individualism and Economic Order*. London: Routledge and Kegan Paul.
- [17] HIPAA. Break glass procedure: Granting emergency access to critical ePHI systems. Retrieved on Jan, 11 2009, 2009.
- [18] J. Jaisingh and J. Rees. Value at risk: A methodology for information security risk assessment. In *In Proceedings of the INFORMS Conference on Information Systems and Technology*, 2001.
- [19] B. Karabacak and I. Sogukpinar. Isram: information security risk analysis method. *Computers & Security*, 24(2):147–159, 2005.

- [20] A. K. Lenstra and T. Voss. Information security risk assessment, aggregation, and mitigation. In H. Wang, J. Pieprzyk, and V. Varadharajan, editors, *ACISP*, volume 3108 of *Lecture Notes in Computer Science*, pages 391–401. Springer, 2004.
- [21] J. J. Longstaff, M. A. Lockyer, and M. G. Thick. A model of accountability, confidentiality and override for healthcare and other applications. In *RBAC '00: Proceedings of the fifth ACM workshop on Role-based access control*, pages 71–76, New York, NY, USA, 2000. ACM.
- [22] G. Magklaras and S. Furnell. Insider threat prediction tool: Evaluating the probability of it misuse. *Computers & Security*, 21(1):62–73, 2002.
- [23] J. A. Miller, M. Fan, S. Wu, I. B. Arpinar, A. P. Sheth, and K. J. Kochut. Security for the METEOR workflow management system. Technical report, UGA-CS-LDIS, University of Georgia, 1999.
- [24] A. Moore, D. Cappelli, and R. F. Trzeciak. The “big picture” of insider IT sabotage across U.S. critical infrastructures. Technical Report CMU/SEI-2008-TR-009, CarnegieMellon, May 2008.
- [25] NIST. Fips 65: Guidelines for automatic data processing risk analysis. Technical report, NIST, 1975.
- [26] NIST. Fips 191: Guideline for the analysis local area network security. Technical report, NIST, 1994.
- [27] T. R. Peltier. *Information security risk analysis*. CRC press, 2005.
- [28] D. Povey. Optimistic security: a new access control paradigm. In *NSPW '99: Proceedings of the 1999 workshop on New security paradigms*, pages 40–45, New York, NY, USA, 2000. ACM.
- [29] M. R. Randazzo, M. Keeney, E. Kowalski, D. Cappelli, and A. Moore. Insider threat study: Illicit cyber activity in the banking and finance sector. Technical Report CMU/SEI-2004-TR-021, CarnegieMellon, June 2005.
- [30] E. Rissanen, B. S. Firozabadi, and M. J. Sergot. Towards a mechanism for discretionary overriding of access control. In B. Christianson, B. Crispo, J. A. Malcolm, and M. Roe, editors, *Security Protocols Workshop*, volume 3957 of *Lecture Notes in Computer Science*, pages 312–319. Springer, 2004.
- [31] L. Røstad and O. Edsberg. A study of access control requirements for healthcare systems based on audit trails from access logs. In *ACSAC*, pages 175–186. IEEE Computer Society, 2006.
- [32] J. H. Saltzer and M. D. Schroeder. The protection of information in computer systems. In *Proceedings of the IEEE 63-9*, 1975.
- [33] R. S. Sandhu, E. J. Coyne, H. L. Feinstein, and C. E. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, 1996.
- [34] E. E. Schultz. A framework for understanding and predicting insider attacks. *Computers & Security*, 21(6):526–531, 2002.
- [35] W. Stallings and L. Brown. *Computer security: principles and practice*. Pearson Prentice Hall, 2008.
- [36] G. Stevens and V. Wulf. A new dimension in access control: studying maintenance engineering across organizational boundaries. In *CSCW '02: Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 196–205, New York, NY, USA, 2002. ACM.
- [37] G. Stoneburner, A. Goguen, and A. Feringa. Risk management guide for information technology systems – NIST special publication 800-30. Technical report, National Institute of Standards and Technology, 2002.
- [38] The CRAMM Manager. Cramm user guide issue 5.1. Technical report, Insight Consulting, 2005.
- [39] R. Willison. Understanding the perpetration of employee computer crime in the organisational context. *Information and Organization*, 16(4):304–324, 2006.
- [40] R. Willison and J. Backhouse. Opportunities for computer crime: considering systems risk from a criminological perspective. *European Journal*, 15(4), 2006.
- [41] B. Wood. An insider threat model for adversary simulation. In R. H. Anderson, T. Bozek, T. Longstaff, W. Meitzler, M. Skroch, and K. Van Wyk, editors, *Research on Mitigating the Insider Threat to Information Systems #2*. RAND, 2000.
- [42] X. Zhao and M. E. Johnson. Access flexibility with escalation and audit. In *WISE 2008: Twentieth Workshop on Information Systems and Economics*, 2008.
- [43] X. Zhao and M. E. Johnson. The value of escalation and incentives in managing information access. In *Managing Information Risk and the Economics of Security*. Springer-Verlag New York, Inc., 2009.