Job Design: Providing Strategic Decision Support for Risk Analysis and Policy Definition

Marco Casassa Mont, Adrian Baldwin, Simon Shiu, Paul Collins

HP Laboratories
HPL-2010-35

Keyword(s):
job design, policies, access rights, SoD, decision support, modelling, simulation, security analytics, policy analytics

Abstract:

Strategic decision makers need to organize their workforce and define policies on how to allocate roles and rights to individuals allowing them to work effectively for the organization, whilst minimizing security risks. Many organizations have a separation of duty matrix specifying certain toxic combinations of access rights that they generally understand present an extreme risk. These matrices do not always contain some of the less understood or smaller risks. The flip side of the rights allocation problem is the need for an organization to keep systems running under various pressures including reducing headcounts. This tension often leads to a practice of providing skilled individuals with wide access rights to many systems. We describe this tension as the Job Design Problem. That is how to manage the trade-offs between allocating roles allowing for flexibility and the possible security impacts. It is not just a matter of technical "role engineering", access right allocation and Identity & Access Management (IAM) provisioning processes. Decision makers need tools that help them understand how to give guidance and set policies associated with role allocations and mechanisms to enable a debate between various stakeholders within the business, IT and Audit concerning the appropriate level of tradeoff and acceptable risk. In this paper, we aim at making progress in this field by presenting an approach and methodology to provide strategic decision support capabilities for the definition and assessment of policies in the context of Job Design. We focus on a problem provided by an IT department within a large organization, where employees (primarily IT admins and IT support staff) operate on sensitive and critical business systems and services. In this context, security risks are a major concern and need to be fully understood. Depending on the motivations and skills of the workforce, accidental or deliberate misuses of access rights and capabilities might take place and have huge economical and reputational consequences for the organizations. The decision makers (e.g. CIOs, CISOs) need to understand the implications and trade-offs of making job design decisions as wells as investing in additional/complementary controls, such as monitoring/auditing systems, IAM solutions, education or vetting/clearance programs. We describe a decision support solution based on modeling and simulation, to provide this kind of policy-decision support. This is work in progress. We present our current results and next steps.

External Posting Date: March 6, 2010 [Fulltext] Approved for External Publication
Internal Posting Date: March 6, 2010 [Fulltext]

Job Design:
Providing Strategic Decision Support for Risk Analysis and Policy Definition

Marco Casassa Mont, Adrian Baldwin, Simon Shiu
System Security Labs Hewlett-Packard labs
Bristol, United Kingdom
marco.casassa-mont@hp.com, adrian.baldwin@hp.com,
simon.shiu@hp.com

Paul Collins
Vistorm
Warrington, United Kingdom
paul.collins@vistorm.com

***Abstract*— Strategic decision makers need to organize their workforce and define policies on how to allocate roles and rights to individuals allowing them to work effectively for the organization, whilst minimizing security risks. Many organizations have a separation of duty matrix specifying certain toxic combinations of access rights that they generally understand present an extreme risk. These matrices do not always contain some of the less understood or smaller risks. The flip side of the rights allocation problem is the need for an organization to keep systems running under various pressures including reducing headcounts. This tension often leads to a practice of providing skilled individuals with wide access rights to many systems. We describe this tension as the *Job Design Problem*. That is how to manage the trade-offs between allocating roles allowing for flexibility and the possible security impacts. It is not just a matter of technical "role engineering", access right allocation and Identity & Access Management (IAM) provisioning processes. Decision makers need tools that help them understand how to give guidance and set policies associated with role allocations and mechanisms to enable a debate between various stakeholders within the business, IT and Audit concerning the appropriate level of tradeoff and acceptable risk. In this paper, we aim at making progress in this field by presenting an approach and methodology to provide strategic decision support capabilities for the definition and assessment of policies in the context of Job Design. We focus on a problem provided by an IT department within a large organization, where employees (primarily IT admins and IT support staff) operate on sensitive and critical business systems and services. In this context, security risks are a major concern and need to be fully understood. Depending on the motivations and skills of the workforce, accidental or deliberate misuses of access rights and capabilities might take place and have huge economical and reputational consequences for the organizations. The decision makers (e.g. CIOs, CISOs) need to understand the implications and trade-offs of making job design decisions as wells as investing in additional/complementary controls, such as monitoring/auditing systems, IAM solutions, education or vetting/clearance programs. We describe a decision support solution based on modeling and simulation, to provide this kind of policy-decision support. This is work in progress. We present our current results and next steps.**

Keywords: job design, policies, access rights, SoD, decision support, modelling, simulation, security analytics, policy analytics

I. INTRODUCTION

Enterprise workforces are getting more and more dynamic and complex. Reorganizations, economical crisis and cost cutting efforts, require employees to change their roles or carry out multiple tasks and activities; some of these activities might as well be outsourced to third parties.

In particular, The IT workforce has been affected by these changes. As a side-effect of headcount reductions and changes, employees might be de-motivated or disgruntled or just not have the competences and skills to safely perform their jobs. On the other hand, IT admins and technical staff often need to do sensitive activities on critical business resources. This might resolve in granting the skilled individuals with wide access rights to many systems. In this context, employees can make accidental mistakes with serious consequences or abuse of their access rights to commit frauds and crimes.

It is important, for decision makers, such as Chief-Information Offices and Chief Information Security Offices (CIOs, CISOs), to understand the involved security risks. As a consequence, a major challenge they need to face is how to structure their IT workforce, in terms of who is doing what, and which specific privileges/access rights need to be granted. A variety of conflicting objectives and constraints need to be taken into account, including Separation of Duties (SoD), productivity, limited budgets, actual expertise of the employees.

We refer to this as the *Job Design problem*. In this context, (security) policies need to be formulated along with suitable controls. The Job Design itself can be seen as a possible control that a decision maker can put in place to mitigate the involved risks. Various other alternative and/or complementary controls are potentially available, such as investing in monitoring/audit, IAM solutions, in education and awareness campaigns and vetting and clearance processes.

Decision makers need to identify and understand the involved security risks, explore trade-offs and eventually define guidelines and policies. For example, these policies could dictate that specific roles - e.g. IT security manager and database (DB) manager for critical business services - can only be played by different individuals, to avoid "toxic combinations" of privileges. However, in the real world, the decision maker often needs to trade off security with business and productivity requirements. Policies might dictate "milder"

security constraints and mandate the adoption of monitoring and auditing controls, to mitigate the involved risks.

In general, due to the fuzziness of the IT environment and the involved constraints, defining these policies is not trivial. It is not just a matter of technical "role engineering", access right allocation and Identity & Access Management provisioning processes. Decision makers need tools to enable a debate involving business, IT and audit viewpoints.

In this paper, we aim at making progress in this field. We present an approach and methodology to provide strategic decision support capabilities concerning the job design problem: ultimately we want to provide decision makers with solutions enabling them to explore suitable options and predict the involved risks by keeping into account a variety of aspects, such as people motivations, their skills, business constraints and available controls.

The Job design problem is discussed in more details in Section II. Section III illustrates the proposed methodology and the decision support approach, based on modeling and simulation. Section IV discusses how this methodology has been applied to assess risks and define policies in a customer provided problem. Section V provides the details. This is work in progress. Current results, related work and next steps are illustrated in Section VI.

II. THE JOB DESIGN PROBLEM

The Job Design problem involves a variety of aspects: privileges/rights allocated to people on protected resources; the set of task commonly performed (the term "task" is used in this paper to identify activities carried out on one or more protected systems by using a specific set of rights/privileges); occasional tasks and rights allocated just in specific circumstances; the accumulation of rights during job changes.

In this paper we specifically focus on its implications for an IT workforce. We consider the viewpoint of decision makers and their challenge to understand the involved risks. This is of fundamental importance as it drives the policy decision making process and the adoption of suitable controls.

Let us consider an example, involving an enterprise scenario and a related IT department. The IT department is in charge of managing various IT systems and services, some of them critical to the business, such as: a *Customer Relationship Management (CRM) Service* – which uses a database (DB) containing personal and confidential information; a *Payment Management Service*; a *Testing Environment* used for development and testing purposes, when new applications and/or changes need to be deployed in the CRM service.

In this example, Figure 1 shows how the IT staff and IT admins are currently allocated to different jobs and tasks, based on their skills and competences. This situation might reflect job design decisions made in the past and/or be the result of the evolution, over time, of the IT workforce, to satisfy ever changing business needs and security requirements. In this context, the following security policies might be in place:

- ***Policy1****: There must be separation of duties between IT security and database/data repository administration, for each managed critical service;*
- ***Policy2****: Monitoring and auditing controls must be deployed on managed services and systems to check for security compliance*

Policy 1 specifically aims at preventing employees from leveraging both their DB admin and IT security privileges to tamper with security/audit log files to hide their crimes.

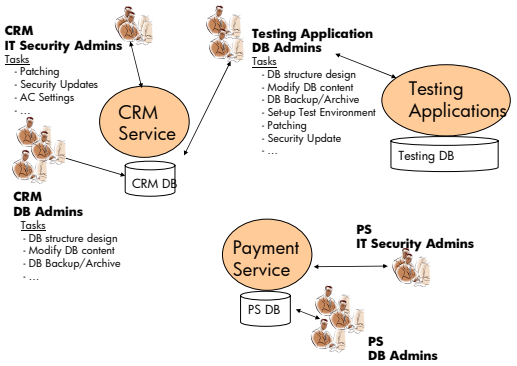

Figure 1. Example of Job Organisation/Job Design

Of course these policies might not be sufficient to deal with all the potential security risks. A disgruntled CRM database admin could still leverage their access rights to copy and resell the content of the database or accidentally damage the organization by destroying its content. A member of the IT staff, in charge of testing applications, could still access the content of the CRM database and deliberately or accidentally leak its confidential content outside the organization.

The overall problem might be more complex that the one described so far, if we keep into account budget and productivity constraints, availability and motivations of skilled personnel and effectiveness of the available controls.

The decision maker (e.g. CISO) might, for example, be facing increasing pressure from the business to reduce costs and streamline the IT support services. Various options might be available, involving reorganization of the IT department and related jobs. For example, Figure 2 shows an alternative way to organize the IT department's job activities, based on a different job design.

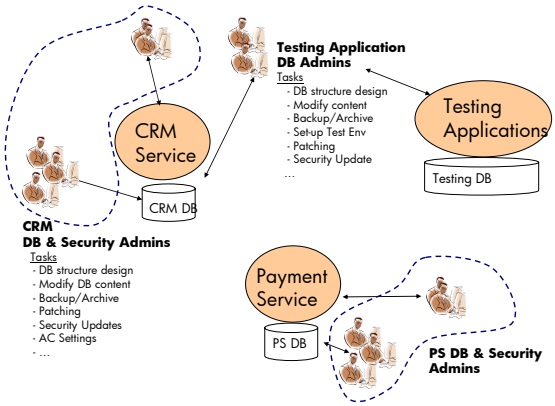

Figure 2. Alternative Job Design

In this case, there is not anymore separation of duties between IT security and database administrators: the two respective teams (for each type of service) are merged. A part of the existing workforce might be laid off, because of these changes and part of the activities might be outsourced – for

example the application testing activities could be carried out by personnel provided by a third party company.

Intuitively this new job design presents additional security risks, due to additional combinations of "toxic privileges" resulting from the elimination of SoD constraints. However, in practice, what is the actual risk? Would this risk be mitigated by introducing additional controls, such as strengthening the monitoring or more qualified/vetted workforce? Which policies should be mandated to effectively deal with the job design changes shown in Figure 2?

We explored this problem driven by a few underlying questions: can we help the decision maker to understand the risk associated to a specific Job Design? If we change the Job Design, can we help understand how risk changes?

There is a unique opportunity to help decision makers to make better/more informed decisions and guide them in defining policies. These decisions cannot be based purely on the analysis of the allocation of access rights: as previously stressed, the overall threat environment, business and security constraints, costs, people's skills and motivations, the effectiveness of additional controls to be deployed, etc. must be taken into account. Decision makers need to consider the various tension points, identify the most suitable trade-offs and eventually make choices to define suitable (security) policies.

This is a complex and risky activity. Decision makers would love to have decision support capabilities that enable them to reason on various Job Design options, explore the implications of their choices, in particular in terms of the involved security risks and understand the mitigations deriving from adopting specific controls. Our work aims at helping them by providing a methodology and decision support solutions to discuss various strategies and controls available to an organisation and define related policies.

III. METHODOLOGY FOR DECISION SUPPORT

This methodology is based on executable mathematical models of the underlying IT systems, processes and threat environments, coupled with methods from risk assessment analysis and empirical data-collection techniques.

Modeling and simulation have been widely used in various fields (e.g. manufacturing processes, environmental and social science) to provide decision support: surveys and data-gathering activities are also used to ground these models. However, their usage in security and IT, coupled with risk assessment methods is relatively new.

Risk assessment (and control management) approaches, such as [7,8,9], provide general purpose criteria to analyse information security risks, by identifying *threat actors*, *potential targets*, *threat vectors* (i.e. potential ways to carry out accidental or criminal attacks) and their impact as well as recommending suitable controls. They usually quantify the involved risks for the worst case scenarios, based on *static* assumptions of threats and their likelihoods. To be effective, they still need to be instantiated in the specific context under analysis.

By coupling risk assessment techniques to modeling and simulation techniques it is possible to explore a wide range of assumptions (about people' behaviors, cause-effects relationships, probabilistic distributions of events and threats, involved processes and workflows, etc.), factor in probabilistic aspects and uncertainty and predict their impact in terms of security risks. Specifically, modelling and simulations techniques are well suited for the Job Design problem because they support the investigation of different hypothesis and the variability of the involved factors, in a context of uncertainty – such as the frequency of accidents and attacks, the dynamic behavior of people, probabilistic effects of security controls, etc. Simulations provide a way to explore them and provide a spectrum of potential outcomes rather than just for the worst case scenario. More details are provided in Sections IV and V.

Recent work by the current authors and others, e.g. [1,2,3,4,5,6] has started to develop this methodology in the IT and security areas and demonstrates its feasibility. Figure 3 provides an overview of the various steps involved in the methodology.

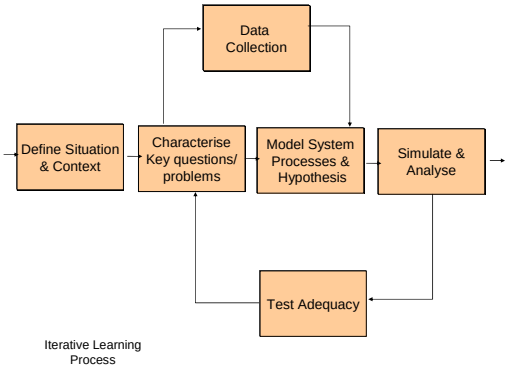

Figure 3. Methodology for Decision Support

After characterizing the aspects involved in the problem, a system modelling phase follows that helps to ground concepts in a specific organisational context; the resulting system model(s) provides predictions of the impact of various choices along with estimates of the outcomes – specifically, the involved risks. This finally helps to identify the most suitable approach and provide guidelines for the definition of policies. Multiple iterations might be required to refine the model and provide effective support to decision makers.

Executable mathematical models can be built to take into account the constraints inherent in the job design problem e.g. people's behaviors, their skills and motivations, tasks to be carried out on specific systems, cause-effects relationships that are at the base of accidents or deliberate attacks, their impact on the organization, mitigation introduced by controls, etc.

The behavior of the model can be simulated in the presence of a (stochastic) representation of the events of relevance - such as frequency of accidents or attempts of attacks - and across different choices involving associations of access rights to people/roles for specific systems, allowed tasks, etc. Its predictions can then be validated against the preferences of the decision makers and various constraints. The model may then be refined appropriately, as the decision maker's understanding of the appropriate targets and preferences in response to the initial problem may itself be subject to reassessment and refinement.

The goal is to help decision makers in thinking about the problem by predicting what happens when changing some of the assumptions, i.e. what if the staff is more trustable or if

monitoring is more or less effective. Modeling and simulation can be used to explore the sensitivity of these issues and focus the decision on taking into account the most sensitive elements. The outcomes are provided to the decision makers in a way they can derive/instantiate their policies by getting a better understanding of the risk implications.

IV. Application to a Job Design Problem

In this section we describe how the proposed methodology has been applied in a realistic setting, specifically a problem provided by the security function of a large IT department. This IT environment involved:

- 140 employees including members of the IT staff and IT admins;
- Critical business systems and services, including a *CRM Service* and a *Payment Service* (PS) and associated *databases/data repositories*;
- An *Application Testing Environment* to test various applications before their deployment in the CRM service.

In this context, the decision maker (CISO) was interested in exploring the implications of reorganizing their IT department by adopting different job designs, different types of controls and get support to define suitable policies. A wide variety of aspects are potentially involved: costs, productivity, security risks, etc. In this paper we focus on the initial exploration we carried out to understand the involved security risks and how to effectively provide decision support to define policies.

A. Concepts and Terminology

The first step consisted in identifying the core concepts involved in the Job Design problem and a suitable terminology to describe them. Figure 4 provides and overview of the various involved concepts:

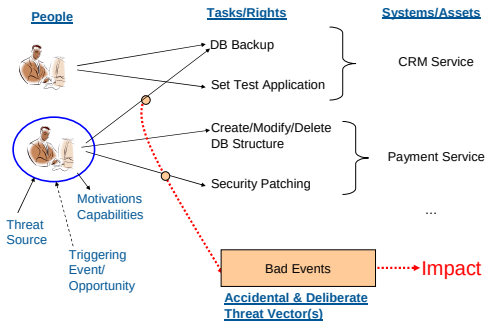

Figure 4. Job Design: Concepts and Terminology

We adopted and extended concepts and terminology used in risk assessment methodologies, e.g., [7,8,9].

People/Agents: these are the member of the IT staff and IT admins. They are characterized by a variety of attributes, including their qualifications, skills and motivations to attack/damage the organization (e.g., by being disgruntled, having financial problems, etc). They have access rights and privileges on protected systems/resources and carry out tasks on these resources. They are the major source of threats for the organizations as they can use/misuse their privileges to accidentally or deliberately damage the organization.

Threat Sources: these could be the involved people (IT staff and admin) and/or additional sources of threat, such as external people/organizations (e.g. third parties bribing employees to access industrial secrets, etc.).

Systems/Assets: there are the valuable resources of the organization that could be targeted in attacks or accidents. In the IT department under analysis these were the CRM and PS services along with their databases.

Tasks: these are activities carried out by people on specific systems. For example patching and upgrading on an IT system, backing-up, restoring or modifying the content/schema of a database, etc. Specific sets of rights and privileges are required.

Triggering Events/Opportunities: these are the events that trigger accidental or criminal/deliberate attacks against the organization. For example, the need, for an IT admin, to pay gambling debts or the fact he/she is disgruntled or doesn't have suitable skills.

Threat Vectors: these are the mechanisms that can be leveraged by people to carry-out *accidental or deliberate* attacks (Bad Events) against an organization. One or more tasks (and related access rights/privileges) on one or more systems might be required to successfully leverage these vectors. One or more individuals might need to be involved. For example, an IT admin that has both DB administrative rights and system security management rights could copy the content of the database and remove any audit trails from the various log files. Threat vectors have potential *impacts* for the organization.

Impact of Threat Vectors: the impact of a threat vector determines the seriousness of the involved risks, based on the consequences for the organization. In our work, for this specific study, we identified ***5 classes of impact,*** in the **[1-5] range**, with an increasing level of seriousness, ranging from limited or negligible consequences for the organization (e.g. minor configuration problems and availability for IT systems), localized impacts (e.g. loss of access to a business resource/service or localized fraud) up to major losses and reputational impacts for the organization (such as massive leakage of confidential data or systemic frauds).

As shown in Figure 4, the tasks (and related privileges) allocated to people on various systems determine which threat vectors could be potentially exploited by these people and the consequential impact. They characterize the "exposure" of the organization to the potential threats posed by its employees. Different allocation of tasks and privileges to people - i.e. a different job design - might resort in a different exposure and related risks.

B. Definition of the Risk Metrics

The second key step consisted in identifying suitable ways to estimate the risks involved in a given job design. In this work the *risk profile* of the organization is defined by *the distribution of the impact,* in the [1-5] range, *of the various threat vectors* that have actually been leveraged/exploited. An example of this distribution is shown in Figures 5 and 9.

Traditional Risk Assessment approaches define the concept of risk (associated to a specific threat/vulnerability) as: ***Risk = Impact * Likelihood.*** The risk depends on the likelihood of a threat happening and its impact. When these approaches are applied to the job design problem, the likelihood and impact

values for the associated *threat vectors* are (usually) statically determined by considering the worst case scenario.

However, this might not necessarily reflect the reality or provide an accurate picture of the actual risks. The likelihood of a person leveraging a specific threat vector might not be deterministic. It could change over time. It might not depends only on static definitions of attributes of the person, such as their skills, motivations and privileges but also on their behaviors and dynamic level of confidence, the knowledge of the presence of specific controls (e.g. auditing/monitoring), etc.

The proposed definition of a *risk profile* based on the *distributions of the impact of exploited threat vectors* takes into account the fuzziness of the problem. In our work, we leverage *modelling and simulation* to determine this risk profile. For example, models can factor in: the current level of confidence of an employee in carrying out an attack by means of certain threat vectors and describe how it changes over time; the frequency and number of people that perform critical tasks; the dissuasion effect of audit/monitoring controls; the implications of introducing security clearances. Simulations (over an observed period of time) can determine the distribution of the impact based on different assumptions made by the decision maker.

V. Providing Strategic Decision Support By Means of Modelling and Simulation

The key goal of our work was to enable the decision maker (in the IT department) to explore the implications that various job design options and choices have in terms of risks and relate this back to the kind of policies they wanted to explore and get support for – based on allocation of tasks/activities to people and controls to be put in place.

Predictive mathematical approaches are suitable to carry out modelling and simulations to achieve this goal. The adopted modelling approach is based on "predictive system modelling", specifically "discrete-event probabilistic modeling" [10, 11]. Our approach, the mathematical basis of which is presented in [2,3,11,12], views a system as having the following key components: **Environment**: it is treated as a source of events that are incident upon the system of interest according to given probability distributions; **Location:** The components of a system of interest are distributed around a collection of places, which may correspond to geographical or more abstract notions of location; **Resource**: this captures the components of the system that are manipulated by its processes e.g. a system, people, etc.; **Process**: this captures the (operational) dynamics of the system. Processes manipulate resources in order to deliver the system's intended services or outcomes.

The adopted approach provides advantages over plain analytical approaches as it explicitly represents the dynamic dependencies and interactions among the involved entities, processes and decisions. This is of relevance for the Job Design problems where a wide variety of events, processes and human interactions are involved. We used the GNOSIS modelling toolset [13] which implements this framework and supports Monte Carlo-style simulations.

We adopted a multi-stage process. We started with an initial model of the aspects involved in a *job design* and expected to run through various different assumptions by changing the model and getting estimates of the involved risks – e.g. based on heuristics around how accidents happen, from customer statistics or using their expertise and implicit experience. In this context, each choice in terms of task allocations and adopted controls has a *direct impact* on the definition of *related policies*.

We then used the model to compare the risk outcomes associated to two different job design options – based on choices made by the decision maker. The comparison of these outcomes provided the decision maker with additional indications of the consequences of related choices. Figure 5 summarizes the core aspects involved in this approach:

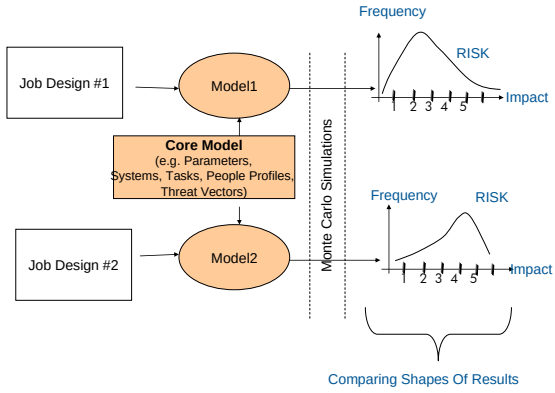

Figure 5. Adopted Modelling and Simulation Approach

A flexible *core model* has been built to represent concepts and processes common to various job designs and control choices. This includes: parameters about events and controls; representation of involved systems, tasks and people profiles (e.g. skills, motivations, clearance, etc.); a way to represent accidental and deliberate threat vectors along with their impacts; parametric processes describing how accidental and deliberate attacks happen, by leveraging threat vectors.

From this core model, *specific models* have been derived to represent the different choices involved each job design and related controls. These models represent: the specific types of jobs defined in the job design; the number of allocated people to each job; their tasks/access rights on the affected systems. Monte Carlo simulations have been used to generate the associated risks profiles, in terms of distributions of the impact of various threat vectors (over a simulate period of time of 10 years), in the [1-5] range of values, as shown in Figure 5.

Ultimately, these distributions enable the decision maker to compare and contrast the implications of making different job design choices and/or adopting different controls and check against their preferences. By exploring various options the decision maker can converge towards an acceptable job design (in terms of the involved risks) and define related policies affecting the allocation of tasks/rights to people and controls to be adopted. More details follow.

A. *Modelling*

Figure 6 provides a high-level view of the specific aspects captured in the core model.

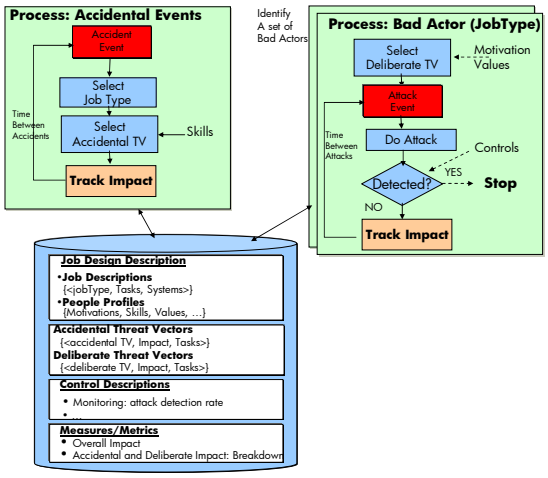

Figure 6. High-level Job Design Model

This model essentially captures the following aspects: *Job Design description; Threat Vector (TV) description; Control description; Measures; Events; Processes*.

The ***Job Design description*** includes: **Set of job types/roles** involved in a job design, the **allocated tasks** and the **systems** that are involved; **Number of involved people** (IT staff and IT admins), **their jobs**, and **their profiles**, inclusive of their skills, motivations (to pursue attacks), clearance level, etc. Specifically, *skill and motivation values are represented in a [1-5] range*, with and increasing value of relevance.

The ***Threat Vector*** description includes: **Accidental Threat Vectors:** each involved accidental threat vector is represented in terms of the tasks necessary to activate it; its **impact** for the organization in the [1-5] range, the impacted system(s); **Deliberate Threat Vectors:** similarly, each involved deliberate threat vector is represented in terms of the tasks necessary to activate it; its impact for the organization in the [1-5] range, *the actual value for the attacker* in the [1-5] range, the impacted system(s).

It is important to notice that the *deliberate threat vectors* are intentionally carried out by employees/agents against the organization. The selection criteria are often based on the value (gain) that the attacker has in doing it, rather than the impact for the organization. As such these two concepts have been explicitly represented in the model.

The ***Control description*** captures: **Set of controls** that has been deployed within their organizations and **associated parameters**. For example, Monitoring/Auditing controls along with their *attack detection rate.*

Measures/Metrics are used during the simulation phase to capture: *the frequency of the exploited threat vectors* for each of the possible 5 classes of impacts, in the [1-5] range; *the break-down details of these values* based on the contribution given by accidental and deliberate threat vectors.

The ***Processes*** explicitly describe how accidental events and deliberate attacks could happen, their consequences, along with the mitigation provided by potential controls deployed within the organization. We made specific assumptions for the specific IT department we analysed which reflect the current availability of data and contextual knowledge. Different assumptions might be true in different contexts.

Specifically, **a process has been modeled to deals with *Accidental Events.*** It is parametric to a *"time between accidents"* distribution which describes how often accidents happen within the organization. Every time an *accident event* happens, this process identifies the more likely "job type" where the accident is going to happen, based on the number of allocated people to jobs and their actual skills. In this model (due to the absence of data coming from the field) we made the assumption that it is more likely to have accidents in those jobs with a large set of people carry out tasks/activities and where people are less skilled. Due to the absence of any statistically significant data from the field, an assumption was also made that all accidental threat vectors have the same probability to happen. Base on this, the process randomly chooses one accidental threat vector (if any) among those that can be activated in that specific job, given the allocated set of tasks. Finally the process tracks the impact of the accidental threat vectors, consistently with the metrics defined in the model.

The model also **explicitly represents the *"bad actors"*** i.e. employees explicitly carrying on attacks against the organisation. The number of bad actors is identified by a parameter that reflects assumptions made by the decision maker. ***Each bad actor is associated to a process***. Each process is instantiated with the job type/role of the bad actor. Due to the absence of empirical data coming from the field, in this model we made the assumption that it is more likely to have a bad actor in those jobs where a large set of people carry out tasks/activities and where people are highly motivated to do criminal activities (e.g. because they are disgruntled or in need for money). This is reflected in the people's profiles defined in the model. The process selects the deliberate threat vector used to carry out the attack (among the ones that can be activated in that job/role – if any) a*mong the ones that bring more value/gain to the bad actor* – i.e. not necessarily the one that has the highest negative impact for the organization. The process explicitly manages the *"time between attacks"* (model parameter). In the context of the IT department under analysis, we made the assumption that this time will decrease when a bad actor successfully carries out an attack, because he/she get more confident. The process captures the risk mitigation factor introduced by the deployed controls. We explicitly explored the impact of deploying a monitoring/auditing control: this is modeled in terms of the *detection rate of attacks*. Each time an attack is carried out the process checks if it has been detected. If it has, the criminal activity (process) stops. Otherwise the process tracks the impact of the deliberate threat vectors, consistently with the metrics defined in the model.

At the simulation time, these processes populate the set of metrics/measures, in the model. This eventually concurs to determine the risk profile distributions shown in Figure 5.

In the IT department analysed in our work, the decision maker wanted to specifically explore the differences - in terms of involved risks – involved in two different job designs, the ones shown in Figures 1 and 2.

In both cases the involved services/systems (assets) were: the *CRM service*, the *Payment Service* and the *Testing Environment system*. In both cases the decision maker also assumed that monitoring/auditing controls were in place, with the same level of investment as well as the same number of employees.

Our analysis identified a set of tasks (along with the associated rights/privileges) carried out by people on the critical systems. This is summarized in Table I:

TABLE I. OVERALL SET OF TASKS FOR EACH SYSTEM/ASSET

System	Tasks
CRM Service	DB structure design; Modify DB content; Backup/Archive/Resume; Patching; Security Updates; AC Settings
Payment Service	DB structure design; Modify DB content; Backup/Archive/Resume; Patching; Security Updates; AC Settings
Test Env.	Set-up Test Environment; DB structure design; Modify SB content; DB Backup/Archive/Resume; Patching; Security Update

Our analysis also identified a set of relevant/core accidental and deliberate threat vectors, based on the above systems/assets and the involved tasks. Due to space limitations, we cannot provide all the details. This information will be available in [14]. Figure 7 shows a subset of these threat vectors, specifically for the CRM service.

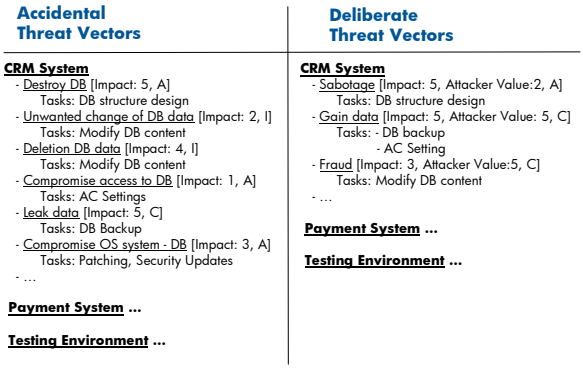

Figure 7. Examples of Accidental and Deliberate Threat Vectors

Figure 7, shows, for example, that a deliberate threat vector could allow the attacker to leak sensitive data from the CRM service by leveraging tasks/rights involving both the back-up of the database and Access Control rights (to eliminate traces of wrong-doing from log files). Figure 7 shows the various attributes associated to the various threat vectors, including levels of impact, types of impact (based on a classification in terms of Confidentiality, Integrity and Availability) and the value for attackers – in a [1-5] range.

The allocations of tasks and people to the two different job designs - shown in Figures 1, 2 - are summarized in Figure 8.

For each job design a model has been instantiated reflecting these various assumptions. Table II summarizes the assumptions made and parameters provided to both models.

TABLE II. PARAMETERS

Events	**Time Between Accidents**: negexp (30 days), **Time for First Attack**: negexp(50 days), **Default Time Between Attacks**: 50 days
Actors	**Number of Employee**: 140, **Percentage Bad Actors**: 3%
Bad Actor Behavior	**Time Between Attacks-Decrease Factor when Undetected:** 1 Day
Control	**Monitoring Control - Attack Detection Rate: 0.001**

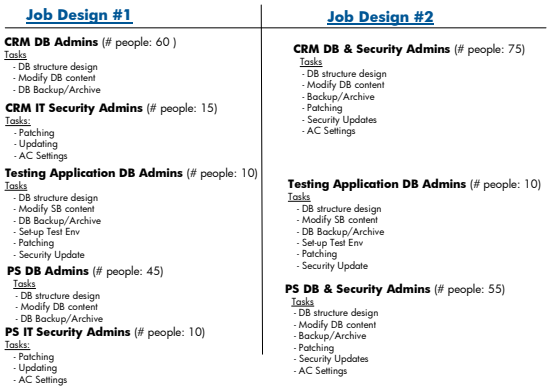

Figure 8. Allocation of tasks and people in two different job designs

B. *Experimental Results*

Monte Carlo simulations have been carried out for each of the two models on a simulated timeframe of 10 years. Each model was run 1000 times to get statistically significant results. Average values have been generated for all the *measures* obtained in output. Figure 9 illustrates the resulting outcomes.

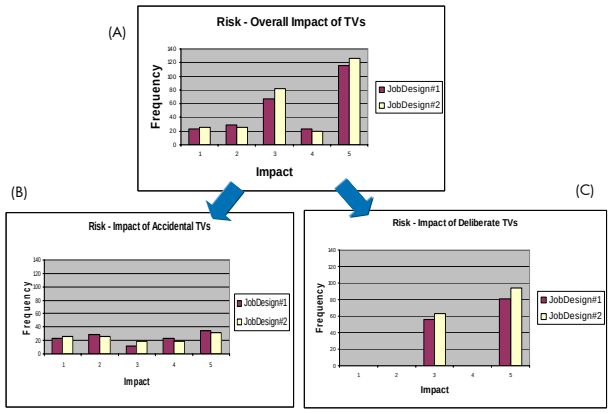

Figure 9. Experimental Outcomes – Risks profiles

Specifically, Figure 9(A) shows the risk profile for the two job designs – by comparing the distributions of the overall impact – in the [1-5] range. As expected from common sense and intuitions, the job design #2 has higher risks - in particular associated to the threat vectors (TVs) which have impact 5, due to the elimination of some of the SoD constraints.

However the results shows there are no major/substantial risk profile differences. This is even more obvious in the breakdown of the measures provided in Figure 9(B) and 9(C) which respectively shows the risk profiles due to the impact of accidental events and deliberate attacks. Figure 9(B) shows that the impact of accidental threat vectors spreads more evenly, due to the wider range of tasks that can be carried out by the people in the involved jobs/roles. In both cases the same assumptions were made in terms of the effectiveness of the monitoring control.

Additional graphs and outcomes can be provided by the simulation, by slicing and dicing the outcomes by types of

jobs/roles, impacted systems/assets and types of impacts. The details are provided in [14].

C. On Decision Support

The above findings have been presented to the decision maker. They have been useful to help them to reason on the actual risks involved in different job designs and drive the process of exploring alternative choices.

In this specific case, the visual comparison of the two risk profiles shows that there are no substantial changes in terms of risks – so it might be reasonable to reorganize the IT department based on Job Design #2. By accepting the involved risks, the decision maker could have used this further evidence to make the job design change and reformulate the involved policies, for example as: ***Policy1****: There must be separation of duties for the IT personnel allocated to different managed critical services;* ***Policy2****: IT security and DB administrative tasks can be carried out by the same team of people, on a specific service* ***Policy3****: Monitoring and auditing controls must be deployed on managed services and systems to check for security compliance.* These policies could then be translated into the operational level in terms of configurations of the systems and allocation of rights and privileges to people.

However, the IT department's decision maker was still interested in further exploring the implications of making different choices, such as reducing the workforce and adopting other controls e.g., investing more on monitoring/auditing or introducing vetting/clearance mechanisms for the employees. All these aspects can be factored in the model(s) and simulations, to determine the new outcomes. This is an important aspect of our methodology, i.e. the possibility to reiterate the process multiple times with the decision maker(s), to explore various trade-off and converge towards suitable solutions that address their risk appetite and preferences.

VI. Discussions and Related Work

The results obtained in this initial work (in the context of the IT department of a large organization) are encouraging: they illustrate how a methodology based on modeling and simulation can effectively help decision makers reason on various job design options and provide predictions that influence the definition of related policies.

This is work in progress. More refined data needs to be gathered from the field, such as the probability of triggering accidental threat vectors for each job and the behaviors of the attackers. This will help to further ground the model. We also need to explore the implications of modeling additional controls of relevance to decision makers, e.g. the vetting process.

In terms of related work, we are not aware of similar R&D activities utilizing modeling and simulations, coupled with risk assessment methods, to provide decision support for the job design problem. As previously discussed, risk assessment methodologies, such as [7,8] are important related work, but they are general purpose, still need to be instantiated to a specific context and their analysis is static and based on worst case scenarios. Similar comments apply to techniques based on attack trees [15]. Despite attack trees help to explore attacks and provide a way to think about security, they could be unmanageable in complex environments, when many potential attack paths and options are available. They cannot easily keep into account dynamic aspects such as people behaviors and controls. Nevertheless, some of the involved concepts can be injected in models – in particular the cause-effect relationships at the base of attacks (or accidents).

VII. Conclusions

This paper focused on the job design problem and the need to provide decision support capabilities to decision makers to understand the implications - in terms of risks - of their choices and trade-offs as well as support the definition of (security) policies. We introduced a methodology based on modeling and simulations to make progress in this space. We discussed how this methodology has been successfully used to explore the problem, in the context of an IT department of a large organisation. Current result and next steps have been presented. This is work in progress.

References

[1] A. Beautement, R. Coles, J. Griffin, C. Ioannidis, B. Monahan, D. Pym, A. Sasse, M. Wonham, Modelling the Human and Technological Costs and Benefits of USB Memory Stick Security. In *Managing Information Risk and the Economics of Security*, Springer, 2009.

[2] Y. Beres, J. Griffin, S. Shiu, M. Heitman, D. Markle, P. Ventura, Analysing the Performance of Security Solutions to Reduce Vulnerability Exposure Windows, ACSAC, 33–42, CA, IEEE, 2008.

[3] M. Collinson, B. Monahan, D. Pym, A Discipline of Mathematical Systems Modelling. Forthcoming monograph, College Publications, London, 2009.

[4] C. Ioannidis, D. Pym, J. Williams, Investments and trade-offs in the economics of information security. To appear, Proc. Financial Cryptography and Data Security, LNCS, Springer, 2009.

[5] A. Baldwin, M. Casassa Mont, B. Monahan, D. Pym, S. Shiu, System Modelling to Support Economic Analysis of Security Investments: A case Study in Identity and Access Management, Trust Economics Workshop and HPL TR HPL-2009-173, 2009

[6] M. Casassa Mont, Y. Beres., D. Pym, S. Shiu, Economics of Identity and Access Management: A Case Study on Enterprise Business Services, HPL Technical Report, HPL-2010-12, 2010

[7] ISO, ISO 27001, Information Security Risk Assessment, http://www.iso.org/, 2005

[8] ISO, ISO 27005, Information Security Risk Management, http://www.iso.org/, 2008

[9] ISACA, Cobit, Control Objectives for Information and realted Technologies, http://www.isaca.org/, 2010

[10] G.S. Fishman, Discrete-Event Simulation: Modelling, Programming and Analysis, Springer-Verlag, 2001

[11] M. Collinson, B. Monahan, D. Pym, Semantics for Structured Systems Modelling and Simulation. To appear, *Proc. Simutools 2010*, ACM.

[12] M. Collinson, B. Monahan, D. Pym, A Logical and Computational Theory of Located Resource. To appear, Journal of Logic and Computation, 2009. Advance Access published on 22 July, 2009.

[13] Gnosis, http://www.hpl.hp.com/research/systems_security/gnosis.html

[14] M. Casassa Mont, A. Baldwin, S. Shiu, A Case Study on Job Design: Providing Strategic Decision Support for Risk Analysis and Policy Definition, to be published as an HPL Technical Report, 2010

[15] B. Schneier, Attack Trees, http://www.schneier.com/paper-attacktrees-ddj-ft.html, 1999