Towards an Analytic Approach to Evaluate Enterprises' Risk Exposure to Social Networks

Anna Squicciarini, Marco Casassa Mont, Sathya Dev Rajasekaran

HP Laboratories
HPL-2009-138

Keyword(s):

Social Networking, Enterprise, Risks, Threats, Modelling, Simulation, Identity Analytics, Security Analytics, Data Leakage

Abstract:

This paper aims at exploring the impact on enterprises of the adoption of Social Networks by employees. It analyses the risks that enterprises could face and suggests a methodology to answer questions, such as: what are the actual risks for an organization, given a specific context? How to assess these risks? What are the most significant approaches that can be taken to mitigate them? What are the financial and organizational implications for an organization in implementing any of the possible approaches?

External Posting Date: June 21, 2009 [Fulltext] Approved for External Publication
Internal Posting Date: June 21, 2009 [Fulltext]

Towards an Analytic Approach to Evaluate Enterprises' Risk Exposure to Social Networks

Anna Squicciarini [1], Marco Casassa Mont [2], Sathya Dev Rajasekaran [1]

[1] The Pennsylvania State University, College of Information Science and Technology,
University Park, PA, USA

[2] Hewlett-Packard Labs, Systems Security Lab, Long Down Avenue,
Stoke Gifford, Bristol, UK

Abstract

This paper aims at exploring the impact on enterprises of the adoption of Social Networks by employees. It analyses the risks that enterprises could face and suggests a methodology to answer questions, such as: what are the actual risks for an organization, given a specific context? How to assess these risks? What are the most significant approaches that can be taken to mitigate them? What are the financial and organizational implications for an organization in implementing any of the possible approaches?

1. Introduction

A Social Network (SN) is a set of people or other social entities, such as organizations, connected by relationships [1]. Social networking sites (SNs) are a type of online community that has grown in popularity over the past several years. They include, among many: Facebook, MySpace, LinkedIn, Twitter, etc. For example, the MySpace SN (www.myspace.com) is ranked tenth in overall web traffic, with over 47 million unique US visitors each month (www.quantcast.com, 2008).

The adoption of SN is getting more and more pervasive, in various contexts. There is an increased adoption of SN also in the context of working activities. The threats and the exposure to security and business risks, arising from careless employees engaging in online communities [2], is now an issue of great relevance for enterprises: employees can disclose not only their personal information but also confidential business data. In this context, enterprises face not only productivity loss due to employee spending time in SNs; a greater concern is the possible threat of information leakage caused by incautious posts or explicit references to private business information. Employees may post information that could negatively impact the company's reputation; write complaints about internal organizational issues or even directly defame the organization they work (or worked) for. The audience of SNs is so broad that besides customers, business competitors, and partners, also hackers may access such information, potentially gaining competitive advantages and causing the targeted enterprise financial losses, both in the short and long term. The risk of attackers exploiting SN data warehouses is on the rise, due also to the tools available to them (e.g. data aggregator and data mining tools).

Decision makers (including CEOs, CTOs, CIOs and CISOs) are now "getting serious about Social Networks" [3], and many companies are proactively studying this phenomenon to understand its possible benefits and risks for the business.

If on the one hand companies can use SNs as a resource to do extensive background checking about potential employees and promote their business, on the other hand, they are now starting to realize the potential risks SN sites expose their enterprises to, and are exploring possible mitigating approaches.

Some initial mitigation measures (controls) include: blocking the incriminated sites [4]; updating security policies to address the admissible SN usage in their "acceptable usage policy" section; introduce new rules and regulations. However, these are reactive approaches to a phenomenon that is not well understood. Many of these approaches are pretty powerless, and prone to fail, if there is no clear understanding of the actual causes of the involved risks and the impact of choices made in an attempt to mitigate them. For example, blocking the sites from office machines helps in reducing the amount of time spent on the SNs, but it is useless for employees working from home, connecting after hours, or using personal devices.

Current operational security (OPSEC) guidance and guidelines about how to have safe online behaviors often resort to the simple suggestion of using common sense. Clearly, this is a not satisfactory solution, and it neither helps prevent nor mitigate possible attacks. Traditional risk assessment methodologies (e.g. ISO 2700x [5]) can help as they provide common sense advices and suggest security-driven assessment processes. However, their recommendations and guidelines need to be refined and contextualized to the SN problem.

We present a conceptual framework for describing and estimating organizations' exposure to threats and risks arising from employee participation in SNs. The aim is to help decision makers to explicitly reason on various aspects related to SN usage, their implications in terms of risk exposure and allow them to explore options to mitigate the possible risks. In order to qualify and potentially quantify the risks implied by SNs from the enterprise perspectives, we propose a new methodology. We begin by providing a *taxonomy of risks* and proceed with *analyzing risks and threats* in the considered domain, by using an approach based on *modeling and simulation*. Models are used to animate and explore different scenarios thereby helping decision makers to understand the implications and consequences of different decisions. Within this context, we use modeling languages, tools [6,7] with strong mathematical foundations to represent the involved processes, users and systems in the enterprise, options and threat environments, thus providing ways to model the implications of possible threats and the impact of different type of controls that the enterprise could potentially employ.

2. Analysis of Threats and Risks for Enterprises

Personal and business data are the key enterprise assets at risk of being exposed in SNs. Employees can both be victims and threat carriers. As victims, they may suffer identity thefts, financial losses and, in extreme cases, lose their job as a result of their unprofessional behavior. With respect to the organization, employees represent threat carriers, since they may expose confidential business data. In addition to more traditional social engineering attacks [8], attackers may indirectly obtain sensitive information from both a single user's profile (on a SN) or by combining different pieces of information, obtained by cross correlating multiple users' profiles belonging to one or more SNs.

From an external attacker's point of view, data can be analyzed, or mined in several ways. We consider the attacker as a malicious person (or organization), possibly a member of the SN, either an insider or an outsider of the targeted company. An attacker can be motivated by a number of reasons, personal and business related, such as financial gain, personal revenge, social activism, etc. Attackers may have a varying knowledge of the domain, and different tools available to reach their goals. Within our context, we abstract attackers' capabilities into three classes: (1) *beginners*, (2) *intermediate* and (3) *experts*. Beginners have limited technical skills, no availability of software or hacking tools to aid their attacks. Intermediate attackers have limited knowledge and availability of software to help their attacks. Expert attackers can run sophisticated attacks, creating new malware, exploit data across multiple SN etc.

Given our interest on enterprise, the focus is on the data exposed by employees rather than the pure "social graph" information [1], i.e. network of links to other people (e.g. friends, colleagues, business partners, etc.).

We identify two related attack approaches, involving the exploitation of data stored in SNs: **Vertical**, which focuses on a specific SN and **Horizontal** where the focus is across multiple SNs. For each of these approaches, the external observer can focus either on a specific person or a group of people. Table 2.1 provides a high level taxonomy of attacks based on how data can be harnessed and processed by an external observer.

More details follow.

Attack Type **Targeted users**	**Vertical (within a Social Networking Site)**	**Horizontal (across Social Networking sites)**
A person	**Attack** Profiling **Risk**: • Identity theft • Enterprise Reputation damage	**Attack** • Extensive profiling (multiple attributes), by linking (when feasible) single profiles of the user on different sites • Making deductions about business aspects based on personal interests/discussion groups **Risk**: • *Personal risk:* identity theft, Blackmailing by using correlated information; • *Business risk: business intelligence* & confidential data leakage with consequential damages for the enterprise
Set of people (e.g. based on affiliation)	**Attack** Cross-checking among profiles for common aspects and inferring implicit personal and business information (e.g. teams/organizational structure, focus areas, research areas, …) **Risk:** High risk of Business intelligence and confidential data leakage with consequential damages (confidentiality issue)	**Attack** • Extensive Cross-checking among profiles for common aspects and inferring implicit personal and business information (e.g. teams/ organizational structure, focus areas, research areas, …) **Risk**: High risk of Business intelligence and confidential data leakage with consequential damages (confidentiality issue)

Table 2.1: High level Attack Taxonomy: Vertical and Horizontal Attacks

Vertical Attacks happen on profiles available in a SN, related to one or more people. For example, a user may post lot of sensitive data in their profile along with job information and data that breaches the security and confidentiality policy of the enterprise (e.g. salary/tax information). Employees often also indicate other potentially sensitive information, such as their full contact information, relationship status, and other social traits.

Some data, disclosed to a SN, might not breach any privacy policy when singularly taken. However, since the amount of data available in a single SN is large, data aggregation is highly possible, potentially revealing sensitive information.

There are two effective approaches to aggregate data coming from different profiles. A fist approach consists of building a *comprehensive user's profile* by collecting attributes disclosed in different profiles of different users who share a same subset of attributes. For example, consider a set of Facebook users who join the group “Employee at the UAHO company in Lewiston”. If any two users reveal the address and the office work number, the work contact information is automatically known for all the individuals belonging to the same group. As discussed in recent studies [9], this is a powerful attack vector in SN, although yet underestimated. Alternatively, the analysis of comments/hypertext/message boards publicly posted can reveal important content. For example, by text-mining certain occurrences of words - like a project name - information about the project’s evolution, status, and outcome may be partially inferred. Furthermore, social applications and widgets can represent powerful tools, to collect additional users’ profile data.

Interesting business information can be inferred when employees update their profiles concurrently. For example, consider a group of users who works in the field of technology, and live in the Philadelphia area. Suppose that after a certain date, many of them gradually start looking out for job opportunities elsewhere. This may indicate an upcoming crisis in the field for that specific area.

Horizontal or Across-SNs Attacks. This analysis is complementary to the vertical analysis. All the observations made before are valid, but now the sources of information are more than one SN site. For example, with regards to aggregated data, an attacker can cross-correlate and complement the attributes of a person’s profile by getting information from their other profiles (e.g. in LinkedIn and Facebook), as long as it is obvious the link between these profiles (e.g. by using registered names, photos, etc.) This approach returns more in-depth results, due to the larger volume of data involved. Notice that this kind of analysis may be resource consuming, especially to link users’ identities across sites. The attacker can however employ some automated crawlers and SN analyzers to lower the complexity of this task.

Besides scripts and other hacking tools, SN provide tools that can be utilized by hackers. For example, Orkut (www.orkut.com) provides the “Polls” application that allows an observer to collect a large number of data on certain topics. A hacker could pose questions aiming at discovering the satisfaction level or the ongoing projects of certain companies. Another potentially effective tool is the “Buzz” tool (http://twitterbuzz.com/) that shows the twitter activity of a company. The company or an employee should own and hold their profile. However, at present, there is no effective control about the veracity of the profiles’ information: a bogus account can be easily created and pretend being an employee of a certain company. Social attacks could be carried out to gain additional information from colleagues. LinkedIn (www.linkedin.com) is also exploitable. It offers, for a limited price, the opportunity of accessing users’ profiles by using fine grained internal searching tools. Hackers could invest less than 25 dollars per month to access users’ profiles (even if they are not part of their networks of friends/colleagues), conduct detailed searches and retrieve related information, references, resumes and read descriptions (when available) of work activities.

3. Risks Assessment and Strategic Decision Support in Enterprises

This section introduces our methodology for assessing risks in enterprises and providing decision makers with decision support tools, to explore the impact of their potential decisions and choices in this space.

We aim at using discrete-event probabilistic modeling and simulation [6,7] to: (1) represent, explain and predict the impact (in terms of data leakage/loss) that employees have on enterprises by exposing information on SNs sites; (2) provide indications about the impact and effectiveness of risk mitigation choices that could be made by decision makers. To achieve this goal the attackers' perspective has to be factored in, by keeping into account their skills and motivations.

3.1 Modeling

Two key basic questions are of relevance, to address key decision makers' concerns:

1. What is the current level of risk and data exposure, for a given organization and a given threat environment?
2. What are the consequences, in terms of risks and costs, of making certain choices, investment or policy decisions? What are the potential best investments, given a limited set of resources?

In general, decision makers have different "levers" (i.e. means of achieving desired outcomes) they can act on to change employees' behaviours, based on available resources and investments. In this context, we considered the following:

- **Education:** investments can be made to educate employees, create awareness about related threats when using them, along with correct behaviours;

- **Monitoring, Awareness and Punishment:** investments could be made in monitoring users' behaviours (using auditing and/or logging), creating awareness of unacceptable behaviours and punishing them, based on collected evidence;

- **Control points:** investments could be made on "control points" (controls), such as software and hardware solutions to monitor, intercept and block data leakages (e.g. by using email interceptors, DRM solutions, black-listing and blocking of web sites).

Implementing a lever is not purely a technical problem, but it has an economic impact too, considering that limited resources are usually available. Ideally, investments in this space should be optimally chosen, to mitigate the risk up to a point where the marginal cost of implementing controls is equal to the value of additional savings from security incidents. In practice, when talking about security it is hard to discuss about *Return of Investments* but rather it is more reasonable to talk about *Value-at-Security Risk* [10]. Here, the value-at-security risk is personal and confidential data that could be leaked and harnessed by third parties, with consequential damages for an organization and its employees. A decision maker needs to understand, at least qualitatively, the impact that employees, in a specific context, have on data leakage. On the other hand, the decision maker needs to assess the impact of remediation choices and the value it can get out of them. Restrictive controls, such as strong forms of access control [11] or monitoring, may be expensive to implement on a large population of users, and in certain domains they might not even be effective as people might find creative ways to bypass them. The involved population of employees may also negatively react to some controls (including a drop of performance) or certain controls may require an adaptation phase, during which the risk will not be mitigated. Finally, based on the specific threats identified, the profiles of involved employees and used SNs, some levers may be more effective than others. For example, for a high skilled, professional workforce, some additional education courses might provide the right level of awareness on how to deal with external SNs, whilst for other employees it might be required to adopt monitoring & punishment controls or control points blocking the access to SNs and the disclosure of data.

Figure 3.1 provides a high-level description of a conceptual model showing the entities and elements of interests.

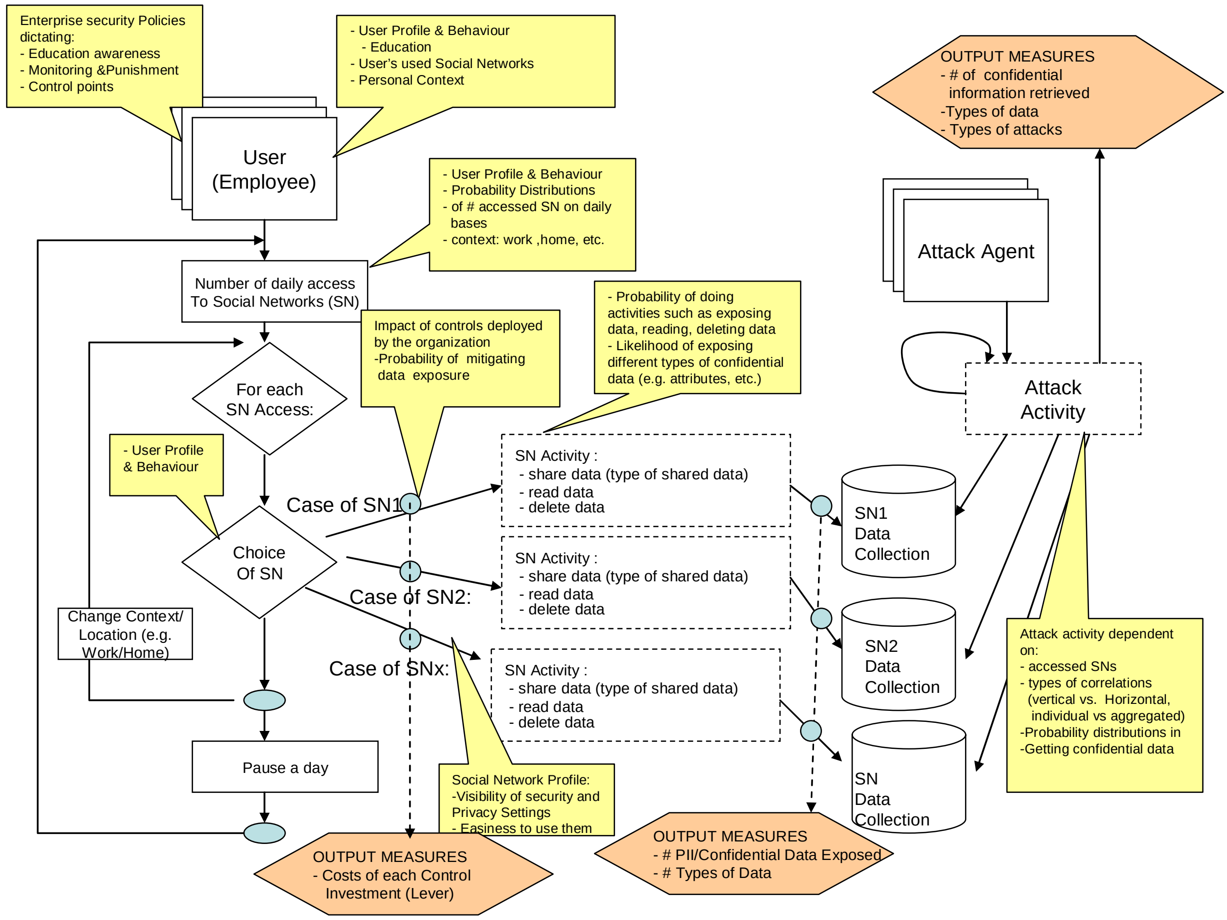
Enterprise security Policies dictating:
- Education awareness
- Monitoring &Punishment
- Control points
- User Profile & Behaviour
- Education
- User's used Social Networks
- Personal Context
User (Employee)
OUTPUT MEASURES
- # of confidential information retrieved
-Types of data
- Types of attacks
- User Profile & Behaviour
- Probability Distributions
- of # accessed SN on daily bases
- context: work ,home, etc.
Attack Agent
Number of daily access To Social Networks (SN)
Impact of controls deployed by the organization
-Probability of mitigating data exposure
- Probability of doing activities such as exposing data, reading, deleting data
- Likelihood of exposing different types of confidential data (e.g. attributes, etc.)
Attack Activity
For each SN Access:
- User Profile & Behaviour
SN Activity :
- share data (type of shared data)
- read data
- delete data
Case of SN1:
SN1 Data Collection
Choice Of SN
Case of SN2:
SN Activity :
- share data (type of shared data)
- read data
- delete data
SN2 Data Collection
Change Context/ Location (e.g. Work/Home)
Case of SNx:
SN Activity :
- share data (type of shared data)
- read data
- delete data
Attack activity dependent on:
- accessed SNs
- types of correlations (vertical vs. Horizontal, individual vs aggregated)
-Probability distributions in
-Getting confidential data
SN Data Collection
Pause a day
Social Network Profile:
-Visibility of security and Privacy Settings
- Easiness to use them
OUTPUT MEASURES
- # PII/Confidential Data Exposed
- # Types of Data
OUTPUT MEASURES
- Costs of each Control Investment (Lever)

Figure 3.1: Conceptual Model for Enterprise Data Leakage Due to Employee Adoption of SNs

This conceptual model factors in the SN activities carried out by employees on a daily basis, both at work and home. These activities include the possibility that personal and business data is leaked to external SN sites. Disclosed data can be exploited by activities carried out by "attack agents" (i.e. external observers, such as hackers and other attackers).

Specifically, a qualitative analysis of risks needs to keep into account the following key aspects:

- **Users** (employees) with different behaviours and attitudes to SNs. In general, users are more or less likely to use SNs and or to disclose information based on skills, level of education and awareness of potential punishments. Each user is modeled as an autonomous agent that performs some of the following activities:
    - Subscribes to a set of SNs and interacts with SNs in different contexts (at work, home, etc.);
    - Accesses SNs a number of times a day, as defined and dictated by a probability distribution. Their attitude to SNs depend on the education, awareness and the place where they operate;
    - Performs a set of activities (based on probability distributions) within a SN, e.g. adding new material, deleting it, reading information.

  The model can measure the number of exposed items, their types, where it has been exposed (type of SN) etc.

- **SN sites**:
    - A SN can be abstracted and modelled as a "data storage & set of services" where users can add information, share, delete or read it. Some degrees of access control and protection might be available;
    - A SN is characterised by the number of subscribes and volume of stored data. External entities might also access and read this information;
    - A SN enforces a defined level of security and privacy controls**.**

- **Attack agents aiming at getting confidential data stored in one or more SN**s:
  The model makes assumptions about the population of attackers that operates to harness and exploit data. Each attack agent is modelled by:
    - A profile that includes motivations and skills in carrying on their actions. These aspects are modelled with probability distributions that qualify the frequency of attacks (e.g. based on motivation) and likelihood in succeeding;
    - The likelihood they could gain access to one or more SNs is also based on skills and the available tools and types of attacks.

  The model measures the amount, type and value of data accessed by the attacker based on their success rate.

  To effectively calculate the risk level and the actual impact of data leakage, it is important to make assumptions on the "threats" i.e. the number of attackers, their skills and their motivations. For the sake of simplicity, the model illustrated in this paper considers three types of "threats", Low, Medium and High, each characterised by a different instantiation of the above parameters.

- **Different "levers" enterprise decision makers can act on and their impact.** As anticipated, these levers (operating in the [0,3] range – i.e. adoption: *none, low, medium, high*) include: *control points* (CP_L), *education* (ED_L), *monitoring and punishment* (MP_L) levers. The impact of these levers is modelled using probability distributions that describe the likelihood that certain "data leakage events" are actually not going to happen as mitigated by related investments.

3.2 Current Model and Experimental Evaluation

In order to provide decision support to decision makers, an "instance model" has been fully implemented, by using the AnyLogic modeling and simulation framework [6]: it is based on the macro-model illustrated in Figure 3.1. The power of an approach based on modeling and simulation is that it allows decision makers to explore a variety of assumptions and choices (e.g. in investment levers) along with providing explanations and predictions of their impacts, in terms of costs and risk exposure. A Monte Carlo approach is used to obtain statistically significant outcomes from simulations.

The implemented model first must be tuned to ensure that it reflects the current enterprise situation (e.g. in terms of detected data leakage and/or current investments in various levers). Secondly, it is used for "what if" analysis and predictions of relevance to decision makers.

This model is characterized by a set of input parameters, each qualifying profiles and properties of the involved entities. A population of 15000 employees (making use of SNs) has been considered. The simulation time has been tuned on a period of 3 years.

This model keeps track of a set of output measures that qualify the outcomes:

- Amount of data that is disclosed by users (data leakage);
- Types (value) of data leaked;
- Actual amount of data that has been successfully accessed by attackers (exploited data leakages), based on the various categories of attacks we identified;
- Actual types (value) of data accessed by attackers;
- Investment costs made by the enterprise for each lever the decision maker has decided to act on.

These output measures are basic measures that can be used to carry out a series of analysis and determine the value of metrics that could be of relevance to decision makers. In particular, we identify two types of High-Level Metrics of relevance for decision makers:

- **Overall Investment Cost for Enterprise**

- **Overall Risk for the Enterprise**

The ***"Overall Investment Cost for Enterprise"*** is modeled as the weighted sum of ***"Fixed/Cost Initial Investment"*** and ***"Variable Cost/Maintenance Cost"***.

The **Fixed/Cost Initial Investment** has been modeled as a linear equation:

Fixed Cost = A*CP_F + B*ED_F + C*MP_F

where: CP_F, ED_F and MP_F are 0 if the corresponding levers CP_L, ED_L and MP_L are 0; 1 otherwise. A, B and C are weights (e.g. set to A=50; B=10; C=30);

The **Variable Cost/Maintenance Cost** keeps into account the time duration for which the investments in levers are made and the number of employees in the enterprise. The variable cost has been modeled as:

Variable_Cost(t)= D(t) * CP_L * t + E(t) * ED_L * t + F(t) * MP_L*t

where: CP_L, ED_L and MP_L vary in the [0,3] range and D(t), E(t) and F(t) are variable weights reflecting different impact and costs of levers in different period of times. The cost of levers may change over time due to software updates and license renewal; training/education sessions; upgrade of the monitoring system.

The "**Overall Risk for the Enterprise"** keeps into account measures collected by the model and reflecting the behaviors of employees:

- ***Info_Potentially_ Disclosed****= Personal_info + Professional_info*: The total information that can be potentially disclosed by employees in the SN based on their activities.
- ***Info_Actually_Disclosed*** *= function(CP_L,ED_L,MP_L) * Info_Potentially_ Disclosed***:** the actual amount of information that is disclosed to the public depending on the various investments in the "levers" used by the enterprise to reduce data leakage (actually a function of these levers, combining their efforts).

The "Overall Risk for the Enterprise" deriving from an attacker is calculated as:

Overall Risk Explosure = function(skill_level, motivation) * Info_Actually_Disclosed

where the "skill_level" of the attacker can take the values between [0,2] (i.e. *none, basic, high*), as well as the "motivation" of the attacker, that determines the frequency of attack.

We assume that, users join the SN by initially exposing little or no information in their profiles. Updates are performed on regular basis. We use a random distribution to check if the user is performing an update. The update frequency can increase based on the user's profile. We make the assumption that a highly educated employee is more aware (and less willing) of disclosing personal and confidential information on the SN [12].

As discussed, the risk associated with an enterprise not only depends on the amount of information available to the SN but also on the characteristics of the attackers. As anticipated, we consider the skill/capability and motivation to characterize an attacker. Attackers can carry out both vertical and horizontal attacks (see Section 2). In our model, the attacker chooses to carry out a specific type of attack based on a random distribution which is driven by their skills and motivations.

Enterprise investments have been modeled around the three levers they can act on - in the [0,3] range - the higher the value of the lever, the more effective it is:

- **Control Point Investment:** We assume the initial installation cost to be a specific amount (e.g. $100000). Fixed and variable costs are kept into account.
- **Education/Training:** When the organization spends on training users are more aware of the implications of using SNs and they will tend to expose less information than they would have done, without training. The reduction in information for a user is based on a triangular distribution, *triangular (min, max, mode)*. The minimum, maximum and the mode values keep increasing as the training lever increases. The higher the training lever, the lower the amount of information exposed and hence lower the risk associated with the organization. Costs increase by years. This is because more people would be involved and will require training and/or need retrain.
- **Monitoring and Punishment:** The higher the value of the investments in this area the more severe the detection and punishment is. If the organization finds that the user has exposed information, it would restrict access of this user to the SN. Thus the user will not be able to perform any updates from the enterprise. The user will however have access to the SN from outside the enterprise. We use a time variable which indicates the place of access of the SN by the user. If the organization has a high degree of monitoring, it will make sure that a higher degree of data leakage will be detected and people punished. This would potentially discourage other people from carrying on similar activities. This is modeled by influencing the probability that a person uses SN, after other colleagues have been punished. The involved costs increase over time. The monitoring system will need to be upgraded if the user has found a way to get past it.

We carried out a set of experiments providing an evaluation of the “threat environment” (low, medium and high), to enable decision makers exploring the impact of different investments for the three different “levers”. A few examples of outcomes of these experiments are provided in the remaining part of this section, to illustrate the type of analysis and results that can be provided to decision makers.

In a first set of experiments we (arbitrarily) fixed the attacker skill level to 0 and the motivation level to 2 to explore the impact that different decision maker’s choices - in terms of investment “levers” - have on the **risk exposure** and the **cost** for the enterprise.

We covered all possible combinations of the three levers, totaling a number of 64 experimental runs. Figure 3.2(a) plots normalised values of risk and cost versus the different combination of levers (X axis) - expressed as “*(control, training, monitoring)*” triples.

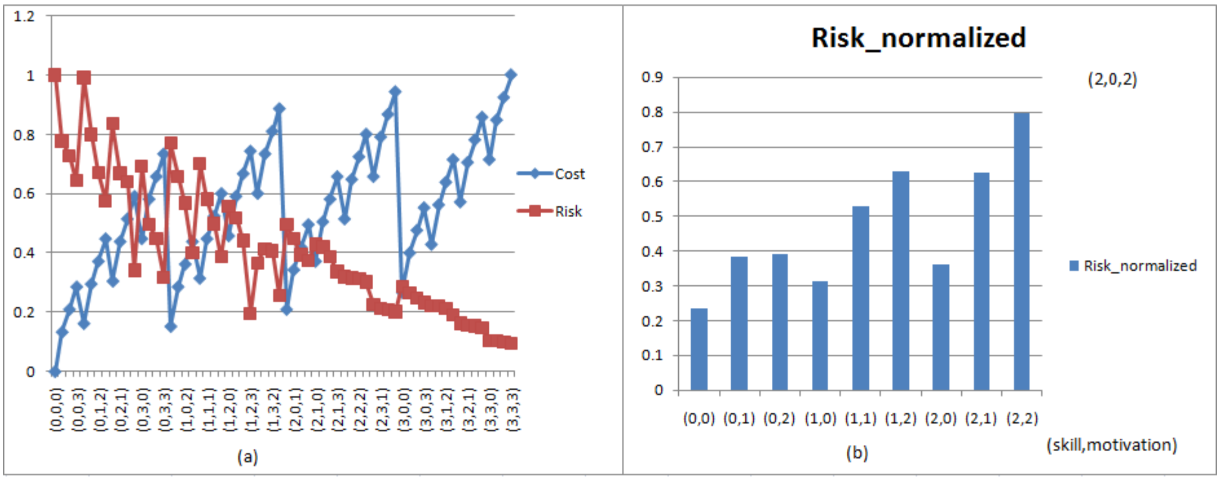

Figure 3.2 (a) Cost and risk simulation outcomes for each combination of the three levers; (b) Variation of risk based on the Motivation and Skill of Attacker

As expected, the cost is minimum and the risk is maximum when the enterprise has decided not to invest in any of the three levers. The cost is maximum and the risk is minimum when the enterprise uses all three levers at its maximum value.

The interesting points are the non obvious ones, in between the two extremes. In the real world, decision makers have to balance (trade-off) the risk exposure with involved costs, due to their limited resources. Thus they would be unlikely to chose the (3,3,3) combination of the three levers, even though this is the most safest combination.

Figure 3.2(a) shows that both the risk and the cost curves resemble a sawtooth curve. Interestingly, in the cost curve there is an enormous dip in the cost, for instance when there is the change of the lever combinations from (0,3,2) to (1,0,0). This is justified by the different cost weights of acting on these levers, over time.

A decision maker, analysing these outcomes might identify a few combinations of interest providing a reasonable trade-off between (normalized values of) risk exposure and costs. These relevant outcomes, resulting from our example, are: (1,1,2), (2,0,2), (3,0,0). Their respective levels of risk and costs are (risk, cost): (0.523, 0.498), (0.419, 0.393), (0.266, 0.286).

Based on the specific assumptions made in the implemented model, for the given level of threats, the (3,0,0) is the most effective combination. The (2,0,2) combination is the second best option.

We carried out a second set of simulations by making different assumptions about the characteristics of an attacker, motivation and skill respectively, to determine how this would impact the overall risk. We considered the case the decision maker uses the (2,0,2) lever combination to minimize data leakage at an efficient cost. The skill and motivation of an attacker both vary in the [0,2] range. A related simulation has been carried out to calculate the risk for all the 9 combinations of the attacker's skill and motivation.

Figure 3.2(b) illustrates the outcomes of this simulation. As we might expect, the risk associated with an enterprise is maximum when the attacker is both highly skilled and highly motivated, in our model this corresponds to skill level 2 and motivation level 2. Based on our model's assumptions, we also observe that the risk of an enterprise depends more on the motivation level than the skill of the attacker. For instance an attacker with skill level 0 and motivation level 2 would cause more risk to the enterprise than an attacker with skill level 1 and motivation level 0. The higher the motivation level of the attacker, the higher is the frequency of attack. Each "attack" time, the attacker would gain a better idea of the types of information available in targeted SNs, for specific people/enterprises and would eventually be able to get relevant information in his successive attempts. This would not be the case of a highly skilled attacker with low motivation, where the attacker would try to get the information in only a few attempts.

In summary, the analysis of our experimental evaluation leads to the following conclusions:

- The risk that organizations face is highly affected by the number of employees that are active on SNs, their behaviors and profiles;
- Given the amount of data exposed in SNs and the lack of built-in security mechanisms protecting enterprises from data leakages, an attack can potentially succeed, as long as the attacker is well-motivated, determined and enough skilled to bypass the authentication and privacy controls;
- Users' education and awareness plays a key role: the more the users are aware about the potential threats associated with data disclosure, the more reluctant they will be in revealing private information and in falling victims of attacks;
- Several alternative solutions exist to mitigate the risk. However, none of these is by itself exhaustive. A combination of "levers" of different nature to act on is desirable to mitigate as much as possible the loss magnitude – but at a cost for the enterprise. The most suitable combination, for a given context, can be determined by using our predictive modeling and simulation approach.

4. Concluding Remarks

In this paper we analyzed the impact on enterprises of the adoption of SNs by employees. We provided an initial taxonomy to classify threats and related risks and used this as a foundation to explicitly assess (1) *the impact of involved risks* and (2) t*he impact of decision makers' choices in this space*.

We proposed a methodology to help enterprise decision makers understanding and assessing the risk involved with SN usage among the enterprise employees. Standards such as ISO 2700x [5], CoBit [13], ITIL [14], etc. describe best practices and methodologies respectively in terms of information security management and risk management, IT governance and service management. Decision makers still need to interpret and instantiate them in their specific operational environments. We use them as drivers and references, but our work further grounds the reasoning to specific contexts and related needs, along with predicting the impacts of specific choices.

We illustrated how modeling and simulation can be effectively be used to help enterprise decision makers to explore the implications of employees using SNs and the impact of their investment choices on risks. In general, models can be used both for quantitative and qualitative analysis, depending on the empirical information available. We argue that both are relevant to allows decision makers to make informed decisions. This work is carried out as a component of the Security and Identity Analytics Project [15], in collaboration with HP Labs.

References

[1] Bonneau, J., Anderson, J., Anderson, R., & Stajano, F. Eight friends are enough: Social graph approximation via public listings. In *Proceedings of Second ACM Workshop on Social Network Systems*. 2009

[2] Computer Weekly, Policies needed to limit social networking risks, says KPMG, http://www.computerweekly.com/Articles/2008/01/10/228852/policies-needed-to-limit-social-networking-risk-says.htm, 2008. Accessed May 31,2009

[3] O. Ross, CIOs getting serious about social networking , ZdNet Asia March 24th 2009 http://www.zdnetasia.com/techguide/security/0,39044901,62051415,00.htm. Accessed May 31,2009

[4] SC Magazine, Companies encouraged to restrict social networking access, http://www.scmagazineuk.com/Companies-encouraged-to-restrict-social-networking-access/article/128244/ Accessed June 1, 2009

[5] ISO, ISO 27001, Information Security Management, http://www.iso.org/iso/catalogue_detail?csnumber=42103, 2005 Accessed June 1, 2009

[6] Anylogic Multi-Method Simulation Software, http://www.xjtek.com/, 2009

[7] B. Monahan, GNOSIS: HP Labs modeling and simulation framework, Systems Security Lab, 2009

[8] Wikipedia, Social engineering attacks, http://en.wikipedia.org/wiki/Social_engineering_(security), 2009

[9] E. Zhelevam L. Getoor, To Join or not to Join: The Illusion of Privacy in Social Networks with Mixed Public and Private User Profiles ACM World Wide Web Conference. April 2009.

[10] R. Hulthen, Communicating the Economic Value of Security Investments: Value at Security Risk, http://weis2008.econinfosec.org/papers/Hulthen.pdf, WEIS 2008, 2008

[11] B. Carminati, E. Ferrari: Privacy-Aware Collaborative Access Control in Web-Based Social Networks. *In Proceedings of IFIP Data And Application Security Workshop.* 2008

[12] Acquisti A., Gross R. Imagined Communities: Awareness, Information Sharing and Privacy on the Facebook. Lecture Notes in Computer Science 4258, Springer, 36-58, 2006.

[13] ISACA, Cobit, IT Governance, http://www.isaca.org/Template.cfm?Section=COBIT6&Template=/TaggedPage/TaggedPageDisplay.cfm&TPLID=55&ContentID=7981, 2008

[14] ITIL, ITIL IT Infrastructure Library for Service Management, http://www.itil-officialsite.com/home/home.asp, 2008

[15] HP Labs, Identity Analytics project, http://www.hpl.hp.com/personal/Marco_Casassa_Mont/Projects/IdentityAnalytics/IdentityAnalytics.htm, 2009