Using Security Metrics Coupled with Predictive Modeling and Simulation to Assess Security Processes

Yolanta Beres, Marco Casassa Mont, Jonathan Griffin, Simon Shiu

HP Laboratories
HPL-2009-142

Keyword(s):
Causal models, simulation, security processes, security metrics.

Abstract:
It is hard for security practitioners and decision-makers to know what level of protection they are getting from their investments in security, especially when they have invested in a number of technologies and processes which interact and combine together. It is even harder to estimate how well these investments can be expected to protect their organizations in the future as security policies, regulations and the threat environment are constantly changing. In this paper we propose that for measuring the effectiveness of security processes in large organizations, a greater emphasis needs to be put on process-based metrics, in contrast to the more commonly used symptomatic lagging indicators. We show how these process-based metrics can be combined with executable, predictive models, based on a sound mathematical foundation, to both assess organizations' security processes under current conditions and predict how well they are likely to perform in potential future scenarios, which may include changes in working practices, policies or threat levels, or new investments in security. We present two case studies, in the areas of vulnerability threat management, and identity and access management, as significant examples to illustrate how this modeling and simulation-based approach can be used to provide a rich picture of how well existing security processes are protecting the organization and to answer "what-if" questions, such as exploring the effects of a change in security policy or an investment in new security technology. Our approach enables the organization to apply the metrics that are most relevant to its business, and provide a comprehensive view that shows the benefits and losses to the different stakeholders.

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

Using Security Metrics Coupled with Predictive Modeling and Simulation to Assess Security Processes

Yolanta Beres
HP Labs, United Kingdom
yolanta.beres@hp.com

Marco Casassa Mont
HP Labs, United Kingdom
marco.casassa-mont@hp.com

Jonathan Griffin
HP Labs, United Kingdom
jonathan.griffin@hp.com

Simon Shiu
HP Labs, United Kingdom
simon.shiu@hp.com

ABSTRACT

It is hard for security practitioners and decision-makers to know what level of protection they are getting from their investments in security, especially when they have invested in a number of technologies and processes which interact and combine together. It is even harder to estimate how well these investments can be expected to protect their organizations in the future as security policies, regulations and the threat environment are constantly changing. In this paper we propose that for measuring the effectiveness of *security processes* in large organizations, a greater emphasis needs to be put on *process-based metrics*, in contrast to the more commonly used symptomatic lagging indicators. We show how these process-based metrics can be combined with executable, predictive models, based on a sound mathematical foundation, to both assess organizations' security processes under current conditions and predict how well they are likely to perform in potential future scenarios, which may include changes in working practices, policies or threat levels, or new investments in security.

We present two case studies, in the areas of vulnerability threat management, and identity and access management, as significant examples to illustrate how this modeling and simulation-based approach can be used to provide a rich picture of how well existing security processes are protecting the organization and to answer "what-if" questions, such as exploring the effects of a change in security policy or an investment in new security technology. Our approach enables the organization to apply the metrics that are most relevant to its business, and provide a comprehensive view that shows the benefits and losses to the different stakeholders.

Categories and Subject Descriptors

K.6.4 [**System Management**]: Management audit, quality assurance.

General Terms

Measurement, Performance, Economics, Reliability, Experimentation, Security.

Keywords

Causal models, simulation, security processes, security metrics.

1. INTRODUCTION

In complex organizations, security processes are used to protect systems of relevance to the business and mitigate related threats and risks. They consist of control points, systems, specific mechanisms, human interactions, and controlled flows of information. Assessing the effectiveness of these security processes at mitigating the risks is not trivial. Multiple control points might be involved to address threats, and across different parts of the IT stack.

Security metrics are created and monitored as a way to get insights about the performance of these controls and to identify failure points or anomalies. However, often the metrics that end up being collected across organizations are low-level, operational metrics, which are amassed without contextualizing them to the overall security processes. Since metrics are often used to predict future behaviors, based on historical data and trends, and to provide decision support for future security investments to different decision makers (e.g. CIOs, CISOs), inappropriate metrics could badly affect those predictions.

We believe that to identify meaningful metrics it is important to take into account the context of the overall security processes and the systems involved as well as to explore the cause-effect relationships that determine specific outcomes for the selected metrics. This is particularly true if these metrics are to be used for assessing the future performance of security processes and for predicting how they would react when circumstances (e.g. threat environment) change.

In this paper we propose using a causal modeling and simulation based approach for the assessment of security controls or processes, and for the identification and exploration of meaningful process-based metrics. Within the context of a causal model we are able to select metrics that reflect an important property or aspect of a security system or process, and through simulations explore why or whether the selected metrics remain good measures under changing conditions.

In particular, we describe two different case studies, on *vulnerability threat management* and *identity and access management*, where we apply this modeling and simulation-based approach. Within the vulnerability threat management case study we assess the current patch management practices in an organization, and identify and explore outcomes for metrics that are indicative of how well the risk from vulnerabilities being exploited is mitigated. In the identity and access management study we assess the implications of moving from ad hoc manual user provisioning to a more automated user provisioning solution. The metrics that are identified and used here directly relate to the evaluation of the new solutions, from a security risk perspective (how much better the new solutions are at mitigated the risk of unauthorized accounts) and in terms of productivity/cost (how much impact in terms of ongoing costs and productivity loss the new solutions would bring).

Our paper is organized as follows. Section 2 provides additional details about different ranges of security metrics. Section 3 describes our methodology, based on predictive modeling and simulation, that when coupled with security metrics, enables to assess security processes and predict future outcomes by means of “what-if” analysis. Two case studies are described in sections 3.1 and 3.2. Section 4 discusses related work. Finally, section 5 sets out our conclusions and some next steps.

2. RANGES OF SECURITY METRICS

Traditionally, security metrics are meaningful measures that can be collected and reported to show whether security controls are working effectively or where risk is emerging. Since security controls are deployed across multiple layers of the IT stack (network, system, OS, application and service, etc.) the range of security metrics that could be or are being collected across these layers is enormous. Most of these metrics are operational metrics, providing focused insights about “day-to-day” security-related events and incidents.

A current dilemma for many organizations is how to make sense of this large set of unrelated, heterogeneous security metrics defined by different parts of the organization, and often gathered at an operational level, and also how to make the metrics more meaningful and indicative of unmitigated risks and security control gaps. The more useful metrics are often the ones that can provide indications of trends and longer-term phenomena and enable the long-term assessment of security processes. Such metrics can then be used to support strategic security decisions, e.g. in terms of security policy changes and security investments.

IT security auditors separate the different types of metrics into two main categories [1]: (1) process-based metrics; (2) symptomatic lagging indicators. Process-based metrics are used to measure an activity or procedure that is part of a control. Such control activities are typically designed by IT management to prevent errors from being introduced into the system, e.g. granting access restrictions to certain capabilities or mitigating potential risks (e.g. using antivirus controls and patching to avoid exploitation of software vulnerabilities). On the other hand, symptomatic lagging indicators are used to measure the effect of the control activity in the data and detect occurrences of errors that may have already been introduced in the system, e.g. a transaction that was improperly authorized.

Table 1 gives examples of the processes and lagging indicators that are often measured to evaluate security-related controls.

The metrics related to symptomatic lagging indicators are often easier to gather, as they might be automatically enabled as part of an IT system’s audit capability. We observed that these types of metrics are often gathered by security operations, based on data in audit logs. For example, within the account management space administrators often collect metrics on how frequently the logons have been used: in the last month, or last half year, also keeping track of the lists of accounts that have been unused or expired. This type of metric is usually generated straight from the user logon audit logs. However, data in logs is usually out of context, and thus is difficult to interpret and analyze. The account management metrics in this example do not indicate how access was granted to the users logging in or if there is an established process to regularly remove expired or frequently used accounts. Such metrics cannot be used directly for business intelligence or for analytical support of security decisions, in this case to identify if any improvements in account authorization or removal process are required.

The metrics collected within the context of a security process or control (process-based metrics) are usually more meaningful measures, as they better indicate the emerging risk in cases when the process is shown not to be working as intended. Good process-related metrics can also indicate the ability of the current security processes to cope with a changing threat environment. These types of metrics, however, are difficult to gather, as they require more contextual information and additional interpretation and analysis. In the example of account management, the useful measures are related to how user accounts are provided and revoked. This would indicate if the account management process is working correctly to preclude any unknown or unprivileged (based on the job function) users being given access, and thus indicate if the systems are placed in an unacceptable risk state.

Table 1. Examples of Process-based Metrics and Lagging Indicators

Security Control Related Activities	**Process-Based Metrics and Tests**	**Symptomatic Lagging Indicators**
Granting, Modifying and Revoking Access	Authorized and unauthorized account ratio measured as part of repeatable process for granting access Ratio of privileged system accounts restricted to IT users and the number of approvals received for each % of privileges assigned that have been checked against job function % of separation of duty conflicts found among users	• Total number of users • Number of users never logged on • Number of inactive users >60 days • Number of locked users • Number of users with expired accounts
Password Administration	Frequency at which password scanning is performed to check complexity Frequency for periodic password changes	• Password Settings • Number of default or unchanged passwords
Patch Management	How fast the patches are being deployed How many patches remain unapplied after policy deadline Difference in patching timeline among critical and non-critical patches	• Number of systems with missing patches based on patch compliance template

The traditional assessment of security processes is done using a “bottom-up” analysis of available data (e.g. logs, various IT measures, etc.) and empirical rules, by identifying metrics that describe the current status, and then these metrics have to be interpreted by the experts doing the risk assessment of a security process. This type of analysis usually takes a considerable amount of time and is often used to support tactical, short-term, and reactive decisions.

It is often impractical and disruptive to do frequent assessment of security processes or to frequently change the security metrics that have to be gathered. In the next sections we show how a modeling and simulation-based approach can be used for the assessment of security processes.

Using this approach we are able to capture the key aspects of a control or process (such as user account management or patch management process) within a conceptual model, and then use this model to explore through simulations the outcomes for the range of metrics that can be gathered as part of this process, and thus identify the metrics that best indicate how the process is working and how risk is being mitigated. By using stochastic simulations, we can explore the effects of various unknown or difficult to obtain contextual information, such as threat environment characteristics, and with the correct metrics we can predict the effect of a new security control or a change in process will have. Once the metrics have been explored through this approach, the security operations staff can put them into effect in the IT environment. With the help of models of security processes, the results of the measurements can then be interpreted on a continual basis.

3. USING PREDICTIVE MODELING AND SIMULATION FOR SECURITY PROCESSES AND METRICS

Modeling and simulation has long been used in the fields of mechanical, civil, environment, and electronic engineering to study how a specific system or design works and to predict behavior under different conditions. By applying modeling and simulation[1] in the area of IT security the aim is to explore how well security processes are likely to perform in current and potential future scenarios, which may include changes in working practices, policies or threat levels, or new investments in security.

In order to apply these methods, we require a conceptual analysis of the key aspects of the security control or process. This is an iterative process involving the analysis of the security control itself together with the wider context in which it is implemented and it often requires interactions with the key domain experts to obtain relevant information and empirical data.

Models are then built by factoring in the relevant representations of the process, human behaviors, and cause-effect aspects. During this phase we also identify how best to measure the outcomes and select a range of metrics to be collected later.

Based on the model, simulations are then carried out using initial assumptions to derive statistically significant outcomes for the range of selected security metrics. Further simulations can be performed for “what-if” analysis to explore the implications of adopting a different security control or of different assumptions in threat environments. Furthermore economics aspects (taking into account costs, impact on productivity, etc.) can be explicitly modeled to drive the analysis from a business perspective.

In the following sections we describe two case studies, in which this modeling and simulation approach has been applied for the assessment of vulnerability and threat management (VTM), and identity and access management (IAM) processes. For both of these case studies we developed models to capture the internal characteristics of the processes in question, and identified metrics to be measured through simulations.

3.1 VTM Case study

Security controls that are usually deployed to deal with software vulnerabilities and related threats, such as malware and viruses, include patch management, antivirus software, and host intrusion prevention systems, with the main prevention control often being patch deployment. Usually, especially in a large organization, thousands of systems running popular business operating systems such as Windows may potentially require patches to be deployed. Deploying patches across all of these systems in a timely manner is not simple. In addition to the time spent on patch assessment and patch testing, the security operations team often faces restrictions on deploying the patches placed by business requirements in terms of limiting system downtime leading to minimal business disruption.

By concentrating on day-to-day operations and applying missing patches, the security team often loses sight of the overall objective of this process: is the risk of malware exploiting vulnerable systems appropriately mitigated by the current patch deployment practices? How exposed are the systems across the organization?

From what we observed, the metrics and indicators that are currently gathered within the vulnerability management area are mostly compliance-driven, and often do not provide an answer to the above questions. In the infrastructure and data center security audits we have been involved in, the main metrics required to be gathered by the security operations consisted of: number of systems that have missing patches (based on a companywide established template of required patches), % of systems that have no antivirus scanning, and frequency of scans.

Although these metrics might be a good indicator of how well the security operations team are doing their task, they do not show if the overall patch deployment is done in a timely manner to mitigate the frequently changing and fast moving threat environment that includes attacks and malware exploiting the known vulnerabilities. A better indication of the effectiveness of the security process would be to show how exposed (in terms of unpatched systems) an organization would be if the malware exploiting a patchable vulnerability arrived at different points in time. Such a metric can be used directly as a measure of how well the processes set by the security policy mitigate the perceived threats. However, this type of metric cannot easily be gathered from historical data primarily because it requires many more samples of the patching process, and a reasonable estimate of the rate at which malware is occurring. That is where simulation-based approach can be applied to derive and explore the outcomes for the different security metrics based on the current practices in patch management within an organization. This type of approach still requires historical data to construct the models as simulation should reflect reality, but the results are gathered from a much larger population which would in many cases more accurately reflect the possible variation in the input events than that observed in historical data. The approach also helps to explore the outcomes in the face of changes to the threat environment, and so ensure the robustness of the selected metrics.

[1] In our case we are using a specialized simulation-oriented language Demos2k [2,3], which implements a mathematical framework based on the foundations of a synchronous calculus of resources and processes, together with an associated modal logic [4].

3.1.1 The Metrics

For the construction of the system model of the patch management process, we first need to examine the vulnerability time line and identify what risk measures are appropriate. Figure 1 shows a timeline of events in the lifecycle of a typical vulnerability. This timeline is similar to that described by Schneier as the "window of exposure" [6] and also examined by McHugh [7] and Frei, et al [8].

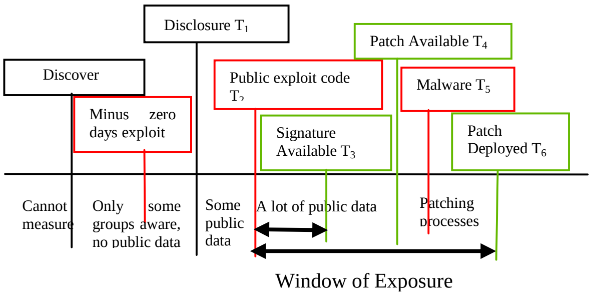

Figure 1. Vulnerability timeline.

Within this timeline the main metric that would indicate how good the standard patching processes are at minimizing the vulnerability exposure is the time taken by security operations team to either (a) deploy the patches or (b) deploy appropriate signatures after the vulnerability was first disclosed. This is represented as the window of exposure in the above timeline. The longer this window, the more the organization is exposed to potential attacks and exploits, as hackers often rush to develop exploits after a vulnerability has been exposed or a patch has been released.

Assuming that we can estimate how quickly these exploits would be developed by the hacker community, the other useful metric to measure is how many systems would be exposed because they were still unpatched at the time that malware hit an organization. This would indicate the magnitude of the risk to an organization in case of malware or attack.

3.1.2 The Model

To assess the specific organization's security processes for vulnerability management and to explore the outcomes for the selected metrics, we constructed a model of these processes, together with characteristics of the external threat environment. We have captured in the model the rates at which vulnerabilities are announced, patches created, and malware emerges based on public threat reports. Conversely, the decisions, timelines and processes within the vulnerability management team in the organization were modeled as discrete event simulations. Figure 2 shows the final model. A detailed description of this model can be found in [16].

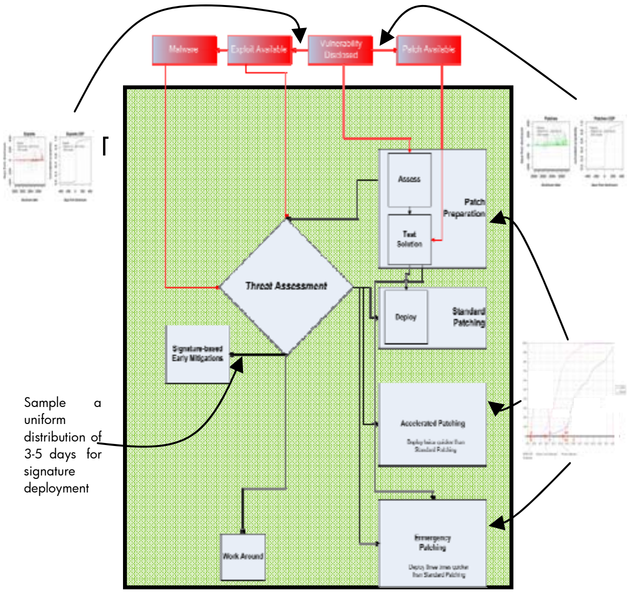
Malware
Exploit Available
Vulnerability Disclosed
Patch Available
Assess
Test Solution
Patch Preparation
Threat Assessment
Deploy
Standard Patching
Signature-based Early Mitigations
Accelerated Patching
Deploy twice quicker than Standard Patching
Emergency Patching
Deploy three times quicker than Standard Patching
Work Around
Sample a uniform distribution of 3-5 days for signature deployment

Figure 2. Vulnerability Threat Management

3.1.3 The Results

The experimental runs with the model were performed through simulations of 100,000 vulnerabilities[2].

We then examined the results for the first metric chosen, the "window of exposure", indicating how long (in terms of days) an organization would be exposed based on its current patch management processes and chosen mitigations. The results are shown in figure 3.

As is visible from this diagram, for around 47% of all simulated cases, it was taking longer than expected (policy-dictated timeline) for risk to be reduced for the organization. This highlights that various delays such as waiting for patches to be released by the software vendor, internal patch assessment and testing, as well as business-related delays, often push the exposure window quite a few days beyond the average, leaving systems exposed for longer than might be expected.

This outcome can then be used directly to make decisions about the improvements necessary in the patch management process to minimize the exposure risk: either by speeding up patch deployment or by finding alternative mitigation mechanisms that can be put into effect earlier. The simulations can then be re-run with the new parameters, and the value of the metric can be examined in the context of the selected process changes to see if the expected timeline is met.

[2] An individual vulnerability had 0.72 chance of having exploit code and 0.97 chance of having a patch, with time delay probability distributions sampled for each event. The time delay distribution for the malware event had a mean of 25 days, and, based on the amount of past emergency patching due to malware appearing in the wild, the probability of malware was set at 0.84 for vulnerabilities where exploit code appeared beforehand.

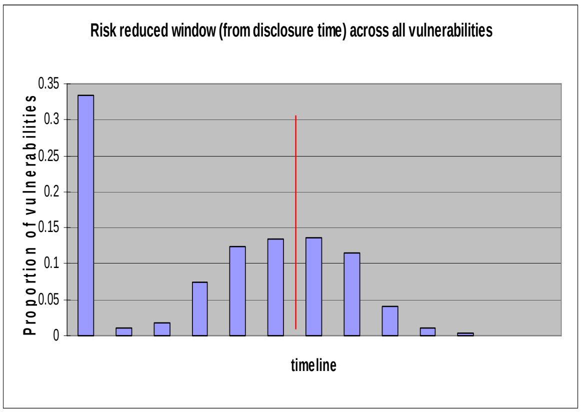

Figure 3. Risk Window Reduction

The outcomes for the second metric, “% of systems exposed at the time of malware” were measured based on some assumptions made about the threat environment. We assumed that malware is being developed for around 60% of all vulnerability instances and that it takes a mean rate of 25 days after public exploit code appears for the malware to be developed and released. Figure 4 shows the results that indicate that in 50% of malware instances most systems would be patched, but there might be a small proportion of cases (20%) where very little of the environment would be patched and where malware could have major impact.

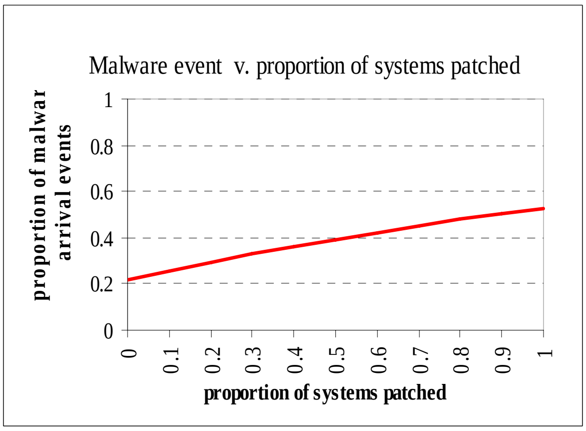

Figure 4. Exposure when malware arrives.

Again the simulations can be re-run with different threat environment conditions, and the changes in the outcomes can be examined to see the increase or decrease in the impact risk.

3.1.4 The Conclusions

Our initial simulations performed for several organizations with slightly different patching processes and the outcomes for the two metrics highlight that in large organizations the vulnerabilities might remain unmitigated for much longer than the expected average, and an intelligent and early threat assessment is crucial for reducing system exposure and impact for highly critical vulnerabilities (the 20% of cases above).

Combining the two metrics provides a good indication of the risk that an organization might be exposed to based on its patching processes. Although the second metric depends on assumptions about the threat environment, and is more appropriate to measure as part of simulations, the first metric, indicating the exposure window, can be implemented as part of other measurements within the actual patch management process of an organization. Rather than collecting metrics just about missing patches, the operations team should track how long it takes to patch the IT environment for each vulnerability that an organization is exposed to and extend this timeline with the time interval from initial vulnerability disclosure.

3.2 IAM Provisioning Case Study

Identity and Access Management (IAM) solutions for enterprises [17] have an impact on multiple aspects of their IT stacks and involve authentication, single sign-on, authorization, auditing, compliance and assurance management, provisioning, data storage, etc.

For the purpose of this case study, we focused on *user account provisioning* solutions. These solutions are used by enterprises to deal with the lifecycle management of user identities and accounts on protected resources. An incorrect or poor provisioning process could give more rights than necessary to users or prevent them from accessing legitimate resources.

We explored the benefits and costs of moving towards a more centralized and automated IAM provisioning process from the current ad hoc manual provisioning processes (e.g. carried out by local system administrators) deployed in many organizations.

From the security perspective there is an implicit assumption that the larger the number of applications managed with the automated account provisioning processes the better risks are mitigated, in particular related to unauthorized user accounts and the least number of compliance issues arising. However, at the same time more automation requires investments (financial costs), and might impact the IT operations team with ongoing operational and support costs, as well as causing unforeseen disruption to productivity of users who have to learn and then work day-to-day within the boundaries of the new processes.

As such, the decision about moving towards a more automated account provisioning process usually requires the input and “buy-in” of several stakeholders in an organization: security experts, who understand the vulnerability of the provisioning process and can articulate the technical consequences; business experts and application/service owners, who understand the criticality of appropriate access to business objectives, and to some extent the business burden the policies create; compliance experts, who are driven by the need to comply with internal guidelines, laws and legislation, pass auditing sessions, etc.; IT Operations experts, who have an understanding of how the IT infrastructure runs along with the related performance, service delivery aspects and costs. Each of these stakeholders usually selects their own metrics and measures for evaluation of the new solution.

The modeling and simulation-based approach is particularly well suited to exploring the implications of adopting and deploying new solutions and processes (in this case the user provisioning process), as it allows experimentation with various assumptions and parameters. We applied this approach to investigate the implications of gradually moving towards more automation of the account provisioning process for the many applications in an organization. To meet the different stakeholder requirements we identified a range of metrics that can be collected during the simulations. The outcomes for the selected metrics can then be used by the different stakeholders to test their own intuitions, share them with others in a coherent and consistent way, and jointly investigate the consequences of a particular investment or policy change in the account provision process.

3.2.1 The Metrics

Different stakeholders care about different metrics measured as part of the IAM provisioning processes.

Traditional *low-level metrics* include: number of correctly configured and misconfigured user accounts; number of hanging accounts (of people who have left the business unit or organization); overall approval time (delays) for provisioning requests; overall configuration/deployment time (delays); number of lost approval and configuration/deployment requests; number of bypassed approval processes. These metrics can be tracked directly from the implemented IAM systems, but they are often only valuable to a subset of the stakeholders (e.g. security administrators and domain experts).

To capture the requirements of all stakeholders involved in evaluating the new account provisioning process, we needed a more wide-ranging set of metrics. Therefore, by carrying out interviews and validating with domain experts, we identified a more comprehensive set of *high-level metrics*, listed below (classified by the relevant stakeholders):

Stakeholder: ***Security/Compliance Officer***

- **Access Accuracy**: the number of correctly configured user accounts, against the overall number of accounts created, including badly configured accounts and hanging accounts;
- **Approval Accuracy**: the number of approved provisioning activities, against the overall provisioning activities, including the unauthorized ones.

Stakeholder: ***Application Owner (Business)***

- **Productivity Cost**: these are the costs, in terms of loss of productivity for employees, due to delays during the approval and configuration/deployment phases of the provisioning process.

Stakeholder: ***IT Operations (IT Budget Holder)***

- **IAM Provisioning Cost**: this is the cost of deploying automated IAM provisioning solutions, for a specified timeframe (fixed and variable costs);
- **Provisioning Effort**: this is the actual number of provisioning transactions carried out by the organization, in a specific timeframe, giving an idea of the effort and involved workload.

It could be argued that some of the high-level metrics above, e.g., productivity cost, are not security metrics *per se,* however they are of direct relevance to the stakeholders, and are essential for decision-making and on-going evaluation of the IAM system by the organization**.**

3.2.2 The Model

To assess the implications of moving towards more and more centralization and automation of the account provisioning process, a detailed model was built that captured stochastically various events such as a user joining, leaving or changing their role. In response to each event, a relevant set of applications was identified where user accounts need to be provisioned/de-provisioned, based on the user's role and profile. For each affected application, either centrally managed or ad-hoc managed, the accounts are provisioned/de-provisioned based on the process steps involved. Figure 5 shows a detail of the process flow triggered by a "User Joining" event.

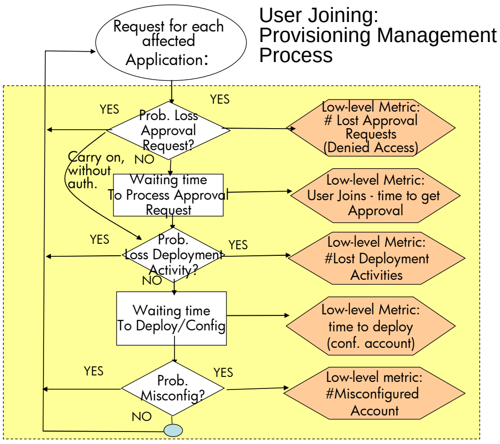
User Joining:
Provisioning Management
Process
Request for each affected Application:
YES
YES
Prob. Loss Approval Request?
Low-level Metric: # Lost Approval Requests (Denied Access)
Carry on, without auth.
NO
Waiting time To Process Approval Request
Low-level Metric: User Joins - time to get Approval
YES
Prob. Loss Deployment Activity?
YES
Low-level Metric: #Lost Deployment Activities
NO
Waiting time To Deploy/Config
Low-level Metric: time to deploy (conf. account)
YES
Prob. Misconfig?
YES
Low-level metric: #Misconfigured Account
NO

Figure 5. IAM provisioning model – Detail on Process and Metrics

This part of the model includes approval and deployment phases, associated delays and failures. The measurements taken as part of this are low-level metrics, such as time taken to get approval, time taken to deploy account on an application, number of misconfigured accounts and so on.

The model includes similar workflows for "User Leaving" and "User Changing Roles" events. Various probabilities for failure and time ranges for delays change depending on whether the application still uses ad-hoc user provisioning or has been moved towards the automated solution. This affects outcomes for various low-level metrics collected across the workflows. More details about our IAM provisioning model can be found [18].

Based on the low-level metrics, as part of the model we also calculate the high-level metrics identified previously. The table in figure 6 shows the formulas used to derive these high-level metrics from measurements taken as part of the various workflows in the account provisioning model.

3.2.3 The Results

In our case study, the organization had 5 core business applications and 100 non-core, lower-priority applications. In the current state, only 2 core applications and 10 non-core applications were provisioned using automated and centralized processes. The organization identified several scenarios ("what-if" cases) for moving towards more automation for both core and non-core applications. These are shown in Figure 7.

Related to these scenarios, experiments were carried out on the model by running simulations over a predefined period of time of 1 year for each case described in Figure 7.

Metrics	Formula	Description
Access Accuracy	1-(w1*UAD+w2*UAM+w3*UAH)/ (UAA)	w1, w2, w3 are relevance weights in the [0,1] range, UAD is the number of denied user accounts, UAM is the number of misconfigured user accounts, UAH is the number of hanging user accounts and UAA is the overall number of user account provisioned (for which either there has been approval or the approval process has been bypassed);
Approval Accuracy	#Approved_Provisioning / (#Approved_Provisioning + # Bypassed_Approvals)	
Productivity Costs	[(join_appr_time+ change_appr_time) + (join_prov_time + change_prov_time)] * *Unit_cost_per_day* + [(#loss_join_appr + #loss_join_prov) + (#loss_change_appr+#loss_change_prov)] **Unit_cost_lost.*	keeps into account loss of productivity due to waiting time (for the approval and deployment phases) and for lost of approval and deployment activities. The impact of these costs are weighted by constants for "unit cost per day" and "unit cost per loss".
IAM Automation Cost	Fixed_Costs + Variable_Costs*Num_IAM_Automated_Apps	Estimated *costs of running* automated IAM provisioning processes, depending of fixed costs (e.g. fixed yearly *fee) and variabl*e costs (e.g. additional license fees depending on the number of provisioned applications)
IAM Effort	#IAM_automated_provisioning_activities	
Ad-hoc Effort	#Ad-Hoc_provisoning_activities	

Figure 6. Definition of High-Level metrics

Experiments	Core Business Applications (5 Apps)	Non Core Business Applications (100 Apps)
CASE #1 – Provisioning CURRENT SITUATION	automation: 2 Apps ad-hoc: 3 Apps	automation: 10 Apps ad-hoc : 90 Apps
CASE #2	automation: 3 Apps ad-hoc : 2 Apps	automation : 40 Apps ad-hoc : 60 Apps
CASE #3	automation: 4 Apps ad-hoc : 1 Apps	automation : 70 Apps ad-hoc : 30 Apps
CASE #4	automation: 5 Apps ad-hoc : 0 Apps	automation: 100 Apps ad-hoc: 0 Apps

Figure 7. Experiments - "What-if" Cases

The simulations produced, as an outcome, statistically significant low-level *metrics* and related high-level *metrics*. Figure 8 illustrates the outcomes for the high-level metrics across the different cases. It shows that accuracy measures are increased by investing more in automation of account provisioning processes. Similarly, productivity costs decrease but provisioning costs increase. For certain investment options (e.g. case 4 - full provisioning automation) the corresponding IAM investment costs are too high, compared to the productivity costs and the actual impact on security.

3.2.4 The Conclusions

The outcomes from the simulations for the low-level metrics and related high-level metrics have been directly used to illustrate the trade-offs of moving towards a more automated account provision process, and provided common ground for discussions between the stakeholders and decision makers involved. In particular, access & approval accuracy metrics (security metrics) were weighed against productivity and IAM provisioning costs (business metrics) to choose the solutions that provided a reasonable trade-off, which in this example was case #3.

In addition, the outcomes for the selected metrics enabled different stakeholders to check the effectiveness of current policies. For example, a policy dictating that "The accuracy of the provisioning process should never be less than 99%" has been demonstrated to be not really cost effective, despite being potentially security effective.

Many of the measurements collected as part of the simulations can be implemented as part of the actual account provisioning process adopted, and can then be used to regularly check if the adopted policies and processes continue to be effective in terms of meeting both security and productivity/cost targets.

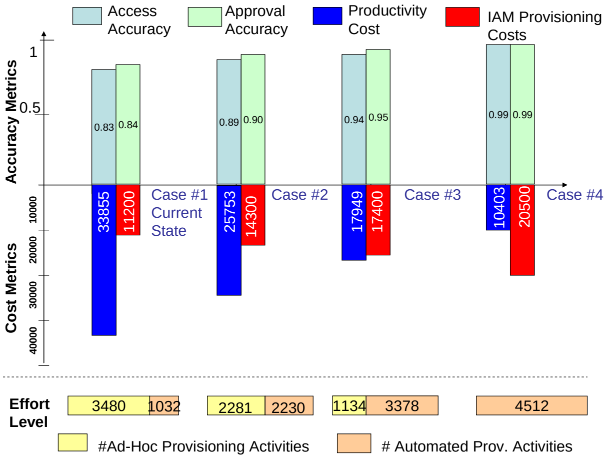

Figure 8. Experiments – Prediction Outcomes for Different "What-if" Cases

4. RELATED WORK

The idea of using analytic and predictive modeling for security has been previously explored [8, 14], but was limited to modeling and simulation of attacks paths and vectors. Wang [13] explores the usage of formal modeling to analyze different security attributes and security flaws in systems. A more detailed usage of predictive models, based on predefined security metrics, is discussed by Gegick and Williams [9], where they use classification and regression trees. Jonsson and Olovsson [10] illustrate how modeling (reliability modeling based on Markov Chains) can be used for preventive security, to predict attacker's behaviors and provide estimates of breaches. Schechter [12] describes the use of regression models in the security domain and discusses the importance of measuring security strength. Deavours et al [11] apply discrete-event predictive modeling systems for the analysis of security properties of the systems. Modeling and simulation has also been extensively used in the system dependability field, with some authors applying some of the methodology to security [15]. Most of this work primarily aims at assessing aspects of security (e.g. vulnerabilities, attack, etc.) by means of different modeling approaches to predict security system or attackers' behavior.

However, none of work referenced above models the high-level security processes (supporting security controls) or examines how security metrics relate to them and analyzes the cause-effect relationships that form the basis of determining these metrics.

This is the core contribution of our paper: we strongly believe that it is necessary to understand the underlying security processes and cause-effect relationships in order to determine the usefulness of security metrics, in providing both long-term predictions (based on historical data) and "what-if" analysis, based on assumptions and hypotheses to be explored.

The metrics related to vulnerability management range from metrics used to evaluate the vulnerabilities themselves [19] to measures of trends of how vulnerabilities have been exploited and how mitigations have been handled by software vendors [7].

McHugh et al [6] were among the first to propose using a window of exposure as a more relevant measure for evaluating how well organizations are doing to minimize their exposure to malware exploiting vulnerabilities. All this work has been used to inform the selection of metrics and the building of the model itself in our vulnerability management case study. There is less in the way of formal related work for security metrics in the IAM space, and primarily concerned with low-level identity management metrics, e.g. [20].

5. CONCLUSIONS

This paper has emphasized the importance of dealing with process-based security metrics vs. more traditional lagging metrics. We have shown how these process-based metrics can be combined with predictive models to provide a variety of benefits to decision-makers, as follows: (1) assess an organization's security processes under current conditions; (2) assess the security metrics themselves; (3) use these metrics to predict how well the security processes are likely to perform in the future, as the threat environment or regulations change; (4) compare alternative security solutions and policies in advance of implementing them; (5) produce consistent metrics at different levels of abstraction suitable for different stakeholders.

We are explicitly modeling the security processes, and the cause-effect relationships between them. A particular strength of our approach is that this causal model provides an excellent grounding for defining the relevant security metrics compared with an approach based purely on correlations of historical data.

We have illustrated how these benefits can be achieved by describing two case studies, involving (a) risk assessment for vulnerability and threat management, and (b) what-if analysis for investments in identity and access management. This work has been further validated by our interactions with HP customers and business groups, and used by them in their IT security decision-making processes.

We plan to build a library of process-based security metrics, with related predictive models, to cover a range of the major IT security areas.

6. REFERENCES

[1] Audit Director Roundtable, "Continuous Control Monitoring: Enabling Rapid Response to Control Breakdowns", Research Findings, February 2005.

[2] Demos2k, http://www.demos2k.org/

[3] G. Birtwistle, Demos – discrete event modelling on Simula. Macmillian, 1979.

[4] D. Pym, B Monahan, "A Structural and Stochastic Modelling Philosophy for Systems Integrity", HP Labs Technical Report Series, HPL-2006-35, http://library.hp.com/techpubs/2006/HPL-2006-35.pdf, 2006

[5] B. Schneier, "Managed Security Monitoring: Closing the Window of Exposure". *Counterpane Internet Security*. http://www.counterpane.com/window.pdf

[6] W. A. Arbaugh, W. L. Fithem, John McHugh. "Windows of Vulnerability: A Case Study Analysis", *IEEE Computer*, 2000.

[7] S. Frei, M. May, U. Fiedler, B. Plattner, ''Large-Scale Vulnerability Analysis", In *Proc. of SIGCOMM'06 Workshops*, September 2006.

[8] A. Atzeni, A. Lioy, "Why to adopt a security metrics? A little survey", In *Proc. of QoP 2005*, 2005.

[9] M. Gegick, L. Williams, "Ranking attack-prone Components with a Predictive Model", ISSRE, 2008

[10] E. Jonsson, A. T. Olovsson, "Quantitative Model of the Security Intrusion Process Based on Attacker Behaviour", *IEEE Transactions on Software Engineering*, Vol. 23, N. 4, 2007

[11] D. D. Deavours, G. Clark, T.Courtney, D. Daly, S. Derisavi, J.M. Doyle, W. H. Sanders, P.G. Webster, "The Mobius framework and its implementation", *IEEE Transactions on Software Engineering*, vol. 28, no. 10, pp. 956–969, Oct. 2002.

[12] S.E. Schechter, "Towards econometric models of the security risk from remote attacks", *IEEE Security & Privacy*,Vol. 3, 2005

[13] A.J.A, Wang, "Information security models and metrics", In *Proc. Of ACM Southeast Regional Conference*, 2005

[14] M., Heidari, "The Role of Modelling and Simulation in Information Security : The Lost Ring", PhD Thesis, http://www.infosecwriters.com/text_resources/pdf/MandS_MHeidari.pdf, 2006

[15] D Nicol, W Sanders, K Rivedi, "Model-based evaluation: from dependability to security. *IEEE Transactions on Dependable and Secure Computing*, Jan-March 2004.

[16] Y Beres, J Griffin, S Shiu, M Heitman, D Markle, P Ventura, "Analysing the Performance of Security Solutions to Reduce Vulnerability Exposure Window", In *Proc. Of Annual Computer Security Applications Conference* (ACSAC 08), December 2008.

[17] M. Casassa Mont, P. Bramhall and J. Pato, "On Adaptive Identity Management: The Next Generation of Identity Management Technologies", HPL-2003-149, 2003

[18] M. Casassa Mont, A. Baldwin, S. Shiu, "Identity Analytics – User Provisioning Case Study: Using Modelling and Simulation for Policy Decision Support in Identity Management", HP Labs Technical Report HPL-2009-57, 2009

[19] Peter Mell, Sasha Romanosky, "A Complete Guide to the Common Vulnerability Scoring System Version 2.0", http://www.first.org/cvss/cvss-guide.pdf, June 2007.

[20] A. Sodbury, "IAM Metrics Case Study", Metricon 1.0, 2006