CHALLENGES TO SUSTAINABLE RISK MANAGEMENT

C. Ariel Pinto, Ph.D., Old Dominion University
Ashish Arora, Ph.D., Carnegie Mellon University
Dennis Hall, Lawrence Berkeley National Laboratory
Edward Schmitz, Ph.D., Office of the Director - NMCI

Abstract

This paper summarizes the intermediate lessons learned from the analyses of the risk management problems in three technological endeavors. These problems are: the absence of a structure for rewarding successful project risk management; the need for an ever-more accurate economic measure of risk; and the difficulty of transferring risks to contract-bound independent outsourcing entity. This paper also describes recent advancement towards providing answers to these challenges and future research endeavors in this field.

Introduction

Risk management is an integral part of any project endeavor. Risks can be classified into three general types which coincide with the three primary concerns in project management: risk of delay in schedule; risk of over-spending; and risk of under-performance. Obviously, these three concerns are very much related such that one affects the other. The amount of available resources affects the rate at which the project progresses, and also affects the overall performance.

Current economic, technological and business situations have renewed the importance of risk management. One such situation is the slowdown of the economy resulting to stricter scrutiny of high-budget government and private endeavors, stricter adherence to planned spending and an even more prudent spending for risk management activities. Another situation is the increasing number and complexity of network-based information systems projects. One of the biggest concerns for this kind of projects is the apparent difficulty of protecting such networks against malicious security attacks and the accompanying uncertainty of the risks involved. Another situation is the popularity of relying on independent organizations for services critical to the progress of the project, or here termed as outsourcing. Together, these situations create a challenging environment for continuous and consistent (i.e. sustainable) management of risks. This is especially true for projects that have particularly high technological components such as software development and maintenance, network infrastructure development, and others.

The central purpose of this paper is to highlight several general challenges in managing risks in the face of these recent situations. These challenges are: a) the absence of a structure for rewarding successful risk management; b) the need for an ever-more accurate economic measure of risk; and c) the difficulty of transferring risks to contract-bound independent entity. These challenges will be exemplified through preliminary analyses of three studies involving the software project risk management at the National Aeronautics and Space Administration (NASA), the economic quantification of information security risk at the Lawrence Berkeley National Laboratory (LBNL), and a proposed incentive scheme for outsourced information security services at the Navy Marine Corps Intranet (NMCI). This paper is organized as follows: a background on basic risk analysis and management is provided, followed by exposition of the three case studies. This is followed by conclusion statements and identified future topics of research.

Background

Risk management can be described in terms of two sets of activities: risk assessment and risk mitigation. Risk assessment can be summarized by posing the following questions (Kaplan and Garrick, 1981): what can go wrong, what is the likelihood that it could go wrong, and what are the consequences. After risks have been assessed, the following questions have to be posed for risk mitigation (Haimes, 1998): What can be done, what are the tradeoffs, and what are the impacts on future options.

During assessment, risk can be quantified as a function of consequences and the likelihood of its occurrence. The Society for Risk Analysis (SRA, 2004) describes risk “based on the expected value of the conditional probability of the event occurring times the consequence of the event given that it has occurred.” Consider a particular scenario Y with likelihood p_Y. Furthermore, suppose that this event has consequence X, a random variable taking on values $x_j, 1 \leq j \leq +\infty$, with likelihood function $p_j = P(X = x_j)$. The expected risk of event Y is shown in Equation 1 below.

$$R_Y = \left(\sum_{j=1}^{+\infty} x_j p_j \right) p_Y \qquad (1)$$

During risk mitigation, it is essential that alternative actions be evaluated based on their costs and benefits, where the benefits naturally will be based on the potential for reduction in risks. Accordingly, the risk of any scenario can be reduced by reducing the consequence x_j's associated with the scenario, or reducing the likelihood of occurrence p_j's and p_Y, or both. However, benefit measured in terms of reduction in risk is not the same as benefit measured in terms of profit. This is particularly true when using financial and economic measures where reinvestment is a basic assumption.

There is also difficulty in accurately estimating the consequence and the likelihood functions associated with a risk scenario. This can be due to the lack of priori information of projects that are unique, and are at the forefront of technology, or involve highly reliable systems such that there are very few historical records to which risk assessment can be based.

Rewarding Successful Risk Management

The current economic slowdown has contributed to the need for a more prudent resource allocation, and this includes resources committed to managing risks. However, there can be a less systemic impediment in implementing sustainable risk management. This pertains to the lack of reward structure for successful risk management that is comparable to other activities that competes for every resource allotted for a project.

In October 2003, the Engineering for Complex Systems Program at NASA's Ames Research Center sponsored a workshop entitled "Managing Software Risk at NASA" held at the Software Engineering Institute (SEI) in Pittsburgh with a goal of raising awareness and forming recommendations for a management strategy for dealing with software risk (Clements and Williams, 2003). This workshop underlined fundamental barriers in implementing software risk management at NASA in the project level. Some of these barriers are:

- Absence of measurable incentive to perform software risk management at the project level
- Risk management are often viewed as added constraint on already limited resources for project completion
- Existing NASA documents in implementing software risk management do not provide tools and techniques for contractors

Consequently, contractors do not know what information to gather and provide to NASA to facilitate software risk management. After occurrences of highly publicized failure of comparable undertakings (e.g. Arian 5) attributable to software component failure, the reaction at NASA is naturally to emphasize better software. Almost all project members at NASA are trained in basic principles of risk management. They are also aware of the divide between projects risks and operational risk management. However, most of their knowledge is on hardware risk management, as opposed to software components of a system. In hardware risk management, NASA employs fairly advanced tools and techniques based on Probabilistic Risk Analysis (PRA), Hazard Analysis, and FMEA. On the other hand, software components of projects are often exempt from such analysis for the reason that the aforementioned tools and techniques can not be readily applied to software components. In a system-level risk assessment, software components are often treated to be not a contributor to system failure.

Furthermore, with the better-faster-cheaper goal and the management-by-number culture at the project level, managing software risk is apparently often not the cheaper alternative and the lack of numbers pertaining to its cost and benefits makes it difficult to manage. Even though NASA management has put in place a large number of policies to handle software risk in a systematic fashion, these documents are more suggestive than obligatory, and do not make clear distinction between hardware and software risk. Samples of these documents are NPG 8000.4, and NPG 2810.1. As a result, various projects within NASA have differing level of implementations of software risk management. As an example, Space Shuttle software is viewed to be the most reliable at NASA. A more detailed discussion of risk management at NASA by Heimann (2000) showed that this situation can be expected on other organizations such as the Food and Drug Administration.

The root of this predicament is the lack of understanding of software risk that is comparable with current knowledge on hardware risk. Current effort by project managers to address the problem is to leverage software risk management during the acquisition process. This entails inserting requirement in the project contracts for contractors to document software risk management based on information gathered for other purposes (e.g. for CMM certification). However, there are no plans yet on how to systematically analyze the information. This, again, is partly due to the lack of formalized procedures for transforming information on software risks to meaningful metrics that describes project success.

Several recommendations brought about by the NASA-SEI workshop that are meant to address this predicament are: update policy and guiding documents to include lessons learned from other projects; have a mechanism for contractors to disclose software-related risk indicators used during development; and avoided risks should be matrixed to the overall project management definition of success (Clements and Williams, 2003, 9).

Recognizing the emerging shift from hardware to software reliability, NASA established the High Dependability Computing Project (HDCP) in 2002. The project has several goals, some of which are the development of methods of providing assurance that particular levels of software dependability are guaranteed, and the long-term improvement of scientific and technical understanding of software dependability (HDCP, 2004). As part of this project, it was identified that one of the ways in addressing the apparent lack of reward structure for successful software risk management at the project level is to develop risk-based cost-benefit methods that can be applied to NASA-relevant software processes, addressing both development risks and operational risks.

Economic Measure of Risk

There is possibly no better way to illustrate the complexity of assessing the economic measure of software risk than in information network security. Nowadays, the security of information network is possibly one of the most challenging applications of risk management. Network security professionals are not simply tasked with implementing more security but also with balancing resources among various technologies for added security. There are multitudes of security technologies to choose from and yet if anything is certain it is that no single technology can guarantee total security - each choice involves risks. The problem then becomes similar to that typified at NASA - a search for structured cost-benefit methods to evaluate and compare alternatives in light of prevailing uncertainties.

In September 2000, the Lawrence Berkeley National Laboratory (LBNL) developed a risk assessment technique for ranking security solutions for their information network based on a quasi return-on-investment metric. This technique established that it is significantly less expensive to accept some damage from cyber attacks than to try to completely prevent all damages. This pragmatic approach enabled LBNL to strike a balance between sound investments on security solutions and acceptable risk of information insecurity.

The technique uses a risk management approach integrating risk profile with actual damages and implementation costs to determine costs and benefits of security solutions. Two crucial concepts are proven to be necessary. The first - incident type, refers to the various categories of cyber incidents that LBNL are able to tell apart. An incident is any undesirable event resulting from attacks against the information system. Although there is no generally accepted incident type naming scheme, most organizations track incidents on an annual basis and group them into types such as root compromise, malicious code (e.g. worms such as Slammer), viruses, and can include inappropriate use and spam email. The second crucial concept - bypass rate, is LBNL's appraisal of the weakness of a given security solution expressed as probability that an attack will be able to penetrate the solution. In effect, each security solution has a bypass rate for every incident type. A 100% bypass rate means the security solution does not stop incidents of that type. LBNL then established procedures to gather data based on these two concepts. The data are described in Exhibit 1 below.

Exhibit 1. Data Gathered at LBNL for Managing Security Risk.

Data	Description
Incident damages	Damage sustained by the institution in a given time period for each incident type and can be approximated by assigning an average cost per incident and multiplying by the number of incidents
Implementation costs of security solution	Implementation and/or operating cost for each security solution for a given period of time.
Bypass rate for each pair of incident type and security solution	Bypass rates can be obtained from vendor specifications, or from white-hat type security evaluation for each security solution. They can also be approximated from interviews with the owners and operators of each security solution.

With these data on hand, the technique undergoes three phases: Calculation of net bypass rate for all security solutions; calculation of total damage, incident risk and baseline scenario; and calculation of risk-based ROI (RROI). Details of the phases are shown in Exhibit 2 below and an example is presented in Appendix A.

Exhibit 2. Details of RROI Calculation.

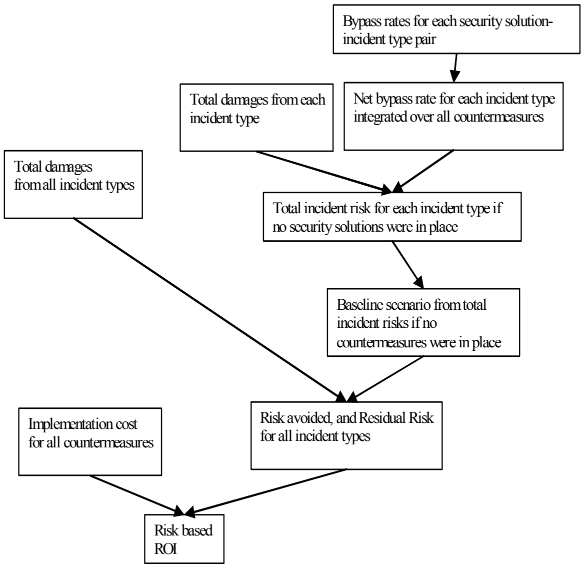
Bypass rates for each security solution-incident type pair
Total damages from each incident type
Net bypass rate for each incident type integrated over all countermeasures
Total damages from all incident types
Total incident risk for each incident type if no security solutions were in place
Baseline scenario from total incident risks if no countermeasures were in place
Implementation cost for all countermeasures
Risk avoided, and Residual Risk for all incident types
Risk based ROI

RROI pertains to the ratio between the net benefit in implementing an IT solution and the cost of implementation. Unlike the conventional notion where return-on-investment measures how effectively resources are used to generate profit, a RROI measures how effectively resources are used to reduce risk. Specifically, a positive RROI means that the degree of risk reduction is greater than the implementation cost, and a greater RROI means more risk reduction per dollar spent in implementation. In essence, a RROI is the ratio between two types of costs: the cost incurred in IT security failure incidents and the cost of thwarting these incidents. Positive RROI does not change the fact that IT security activities are primarily cost centers - those activities that in themselves have negative return on investment but nonetheless provide essential and necessary support for the overall organization.

It is noteworthy that a reduction in risk does not necessarily translate to additional resources which would typically be used for other productive endeavors. In this sense, benefit measured in terms of reduction in risk is not the same as benefit measured in terms of profit. However, the activities leading to the calculation of the RROI provide a security manager a structured cost-benefit method to evaluate and compare IT security solutions in light of prevailing uncertainties.

It is important to note that RROI should be used to guide overall investment in security such that investments should be made until the RROI falls to the minimum rate acceptable to the organization. If, however, one has to choose among alternative security investments, then RROI can prove misleading. Net present value (NPV) is the more robust and consistent alternative measure to ROI when the decision involves choosing among competing solutions. NPV considers the time value of money – the value of a dollar today versus the value of that same dollar in the future, after taking inflation and returns into account. However, the use of NPV poses a burden in requiring more detailed information such as the time when costs and benefits occur. This presents a challenge in security solutions since the occurrence of security failure is highly unpredictable and uncertain. In fact, there are two general challenges in evaluating security solutions: (a) complexity of integrating information on threats, vulnerabilities, attacks, and outcomes, and (b) determining the costs and benefits needed in the analysis. For the framework discussed above, particular challenges are:

Obtaining true costs. Non-cash but otherwise very relevant costs such as lost productivity and opportunity cost of security incidents are often miscalculated primarily due to difficulty in quantifying the actual amount or simply due to lack of enough information. This is particularly true in the valuation of loss of confidentiality and integrity in security breaches. The inherent nature of confidentiality prevents establishing the consequences of security failure, even less putting value on such consequences. However, it is noteworthy that such a challenge also occurs in other settings like physical, health, and environmental risk assessment where human lives are at stake. The implementation cost of the solutions can also be difficult to estimate since many resources, both human and machine, are shared by several solutions during implementation. Double-counting of some costs can also result from vague definitions used in accounting and operation processes. For example, cost due to lost productivity may be difficult to differentiate from cost due to lost revenue. An operations manager may account for work stoppage due to virus attacks as lost productivity, at the same time a financial officer may account for decrease in sales due to the same instance as lost revenue, resulting to possible double-counting.

Estimating bypass rates. The bypass rate, both for existing security system and for the solutions under evaluation can be difficult to estimate due to minimal or non-existent information. Currently, the most reliable sources of this information are intrusion detection experts that have worked closely with the particular solution and have detailed knowledge of the current security system. This is especially true in evaluating new solutions where no actual performance data exist. More recently, there have been developments in using honeypots to directly measure potential frequency of incidents on certain types of networks without using bypass rates. However, bypass rates would still be necessary for calculating residual risk of particular solutions.

Compensating for interaction among solutions. In the example application in Appendix A, the combined effectiveness of the solutions is assumed to be multiplicative, as demonstrated by the calculation of the system-wide bypass rates. However, this simplification may not accurately describe the actual interaction of various solutions implemented concurrently. The architecture of the network and the configuration of particular solutions can result to interaction that may be too complicated to assess, and is beyond the scope of the analysis.

Representing catastrophic losses. A constant challenge in risk assessment is the proper representation of catastrophic incidents. In the example application in Appendix A, it is implied that estimates of costs, consequences, and frequencies are averages or expected values. In this process of averaging out rare but catastrophic events with frequent but inconsequential events, disastrous consequences have the potential to be neglected in the analysis. Though there are tools that deal with this type of events, their demand for detailed information or oversimplifying assumptions often preclude their application in IT security analysis.

Transferring Risks in Outsourcing

The previous section discussed the difficulty in the economic valuation of risk. This difficulty coupled with an apparent efficiency of specialized service providers have resulted to the recent popularity of outsourcing - a managerial and business practice of procuring product or service from another organization under a contract agreement. One good example is the outsourcing of the design and implementation of network security at the Navy Marine Corps Intranet (NMCI) described by Schmitz et al. (2004). NMCI involves the establishment of a single computer network within the Department of Navy replacing over 100 separate organizational intranets throughout the Continental United States and Hawaii.

The general performance of the network is addressed through service levels agreements or SLAs specified in the outsourcing contract. SLAs are quantitative measures which address up to 194 separate metrics of performance for the standard, high-end, and mission-critical parts of the network. Even though an outsourcing contract specifies the minimum level of network performance as measured by the SLAs, there is a need to motivate the contractor to further improve such performance. This is especially true in the fast changing technology of network and information security.

The Carnegie Mellon University in cooperation with the Office of the Director of NMCI formulated and proposed an iterative incentive scheme to motivate the contractor to continuously improve network security at NMCI. The incentive is designed to reward the contractor for demonstrating the ability to defend the network against simulated network attacks as well as actual system performance. Obviously, NMCI has to provide an incentive payment high enough to motivate the contractor towards this outcome but still not overpay for such service. Necessary conditions for a feasible incentive scheme are that NMCI is willing to pay as much as equal to the benefit brought by the increased security level, and that the contractor is capable of implementing such security at a cost less than or equal to the incentive payment (Varian, 1992). This can be expressed as shown in Equation 2, where $c(\mathbf{z})$, $s(\mathbf{z})$, and $x(\mathbf{z})$ are the contractor's cost, the incentive payment, and the benefit to NMCI of a particular level of security $\mathbf{z}$, respectively.

$$c(\mathbf{z}) \leq s(\mathbf{z}) \leq x(\mathbf{z}) \tag{2}$$

Furthermore, the success of outsourcing is also hinged on the outsourcer's knowledge about its own needs and that of the contractor (Bryson and Sullivan, 2002). This is possibly the greatest challenge in formulating an incentive scheme - there is a high degree of uncertainty regarding the benefits and cost of security. Contractors are not required to reveal their costs to NMCI, and for its part, NMCI do not have an accurate valuation of the benefits of incremental improvement in network security and performance. However, the contractor's cost can theoretically be extrapolated through its responses to incentives. One possible approach is to view this predicament as a game where the players are NMCI and the contractor. The first player reveals information and the other player acts based on this revelation. The first player then updates his knowledge and again reveals more information. The process goes on, always in such a way that each player tries to reach their respective goal (Fudenberg and Tirole, 2002). For this particular case, the objective is to define the cost of incentives in a way that relates the benefits of increased security. Given these conditions, the major challenge is to devise an approach to set the target security level. Initially, NMCI can pick a level of incentive $s(\cdot)^{\circ}$ for a given initial budget of I° such that $s(\cdot)^{\circ}$ is less than or equal to I°. From hereon, the superscript $\circ$ pertains to initial values. Note that there is initially not enough information to define the incentive level as a function of the optimal incentive z. The operational issue then becomes how to set a feasible incentive scheme $s(\cdot)^{\circ}$ such that Equation 2 is true for a set of target security level $\mathbf{z}$. It should be noted that since there exists a separate part of the agreement with the contractor that sets minimum SLAs, any increase in security level obtained through an incentive scheme expressed in Equation 2 is at least satisfactory. Also, since the only sources of information available for estimating $s(\cdot)^{\circ}$ and $x(\cdot)^{\circ}$ are the simulated and actual security incidents, it is only but practical to model them as functions of these factors. That is:

$$x(\cdot)^{\circ} = g(z; w_1, w_2, R, A) \qquad (3)$$

where R and A describe simulated and actual security incidents. By the feasibility constraint stated by Varian (1992), Equation 3 becomes

$$s(\cdot)^{\circ} \leq x(\cdot)^{\circ} = g(z; w_1, w_2, R, A) \qquad (4)$$

The following steps describe the proposed iterative incentive scheme.

1. For an initial budget I°, determine an initial incentive scheme $s(z)^{\circ}$ such that $s(\cdot)^{\circ} \leq I^{\circ}$ for budget feasibility, and z greater than the minimum SLAs, and that $w_1 >> w_2$.
2. Observe network performance z against simulated and actual security threats to measure R and A.
3. Extrapolate/Update $x(\cdot)^{\circ}$ and $c(\cdot)^{\circ}$.
4. Modify $s(z)^{\circ}$, w_1, and w_2, and iterate through the steps, such that $w_1 \rightarrow$ nominal level.

Step 1 prescribes setting an initial incentive scheme, possibly based purely on available funds and the minimum SLAs. Step 2 prescribes obtaining information on the assurance level z based on the primary sources of such information: simulated attacks and actual threats. Note that during the initial iterations, it is proposed that most of the incentive payments be made based on network performance against simulated attacks (i.e. $w_1 >> w_2$). This is for the reason that much discovery and exploration is needed at the earlier iterations to extrapolate z. This can only be done through simulated attacks since actual threats may be few and far apart. Step 3 prescribes estimating $x(\cdot)$ based on some function described in Equation 4. Based on z, some extrapolation can be made on the actual cost function $c(\cdot)^{\circ}$. For step 4, as the process progresses and confidence on an accurate $x(\cdot)$ is established, efforts on simulated attacks are marginalized and emphasis is refocused on real attacks (i.e. $w_1 \rightarrow$ nominal level). Also during this progression, the incentive scheme $s(\cdot)$ should eventually be compared to the benefit $x(\cdot)$ and less on the budget I°.

The primary driver of this iterative process is the measured level z. At each period of iteration z can be described as increasing, decreasing, or the same as the previous period. An increasing level could indicate that the incentive scheme is effective. A decreasing or stagnating level of measured z even with an increasing incentive $s(z)$ can denote, among other things, an incentive payment that is too low to be attractive for the contractor. This could also signal that a level have been reached that is technically difficult to increase further.

The convergence of the cost function $c(\cdot)$ towards the benefit level $x(\cdot)$ is another indication of increasing technical difficulty or the effect of decreasing marginal efficiency of improving the security level. Eventually, the cost $c(\mathbf{z})$ for a set of target level $\mathbf{z}$ will approach the budget I, which indicates that NMCI's budget may constrain the attainable security level.

These are some of the issues that must be grappled with in an iterative process of setting incentive budget and information assurance targets:

–How near or far from the optimal incentive scheme as described in the equality scenario for Equation 2 and Equation 4?

–How often should be the iteration?

–How long should be an iteration period?

–How fast can a reliable estimate of the benefit function $x(\cdot)$ be obtained?

–What information needs to be gathered to measure the performance of NMCI against actual threats as suggested in Equation 3?

–What information needs to be gathered to measure the benefit function $x(\cdot)$ as suggested in Equation 4?

For an effective outsourcing of IA services, one has to take into account not only the benefits to the outsourcer from increased level of IA, but also information pertaining to the contractor's costs of achieving given performance levels. The iterative process also provides opportunity for the outsourcer to develop a more mature risk management system through the investigation and documentation of consequences for both simulated and actual threat. Furthermore, the iterative process initiates a closer scrutiny of the existing SLAs: how much are they representative of actual assurance level.

On the other hand, there are challenges to implementing such an iterative process. There are still significant explorations needed to strengthen the functional relationships between the SLAs, the outsourcer's benefits, and contractor's costs. There are also more pragmatic challenges such as those pertaining to the duration and frequency of iteration.

Overall, this paper presented an emerging topic in IA outsourcing and ultimately contributes to the body of knowledge in decision making by offering insights to how outsourcers and contractors interact.

Conclusions

In this paper, we have explored the contributing factors to the absence of a reward process for successful risk management at NASA, LBNL's use of a pseudo ROI based on avoided risks, and a new paradigm in incentive-driven process for outsourced information security services at NMCI. Intermediate lessons learned in these three explorations underline the effects of conceptual complexities of risk and the lack of information in sustainable risk management.

In these three cases, the presence of an acceptably accurate economic measure of risk can contribute to the more sustainable application of risks management. This can provide organizations a way to integrate avoided risk into more traditional concepts of benefits for the purpose of cost accounting and financial evaluation. On the other hand, the inherent lack of information on the actual damage and occurrence of risk scenarios retards it economic valuation. As a result, decision makers need to consider as many criteria as possible for a more robust analysis. This is further pronounced on projects characterized by fast-changing technologies, and uniqueness of application such as those in information technology.

Overall, risk managers are getting a better grasp of investing in risk avoidance, and need to continue exploring unconventional avenues in information gathering and economic valuation.

Acknowledgements and Disclaimer

This work was supported by the Director, Office of Science, Safeguards and Security-Science, U.S. Department of Energy under Contract No. DE-AC03-76SF00098, and by the Software Industry Center at Carnegie Mellon University, and the Office of the Director, Navy Marine Corps Intranet.

While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Appendix A - Application of RROI

The CIO of company X, a medium-sized data service company is trying to evaluate several components of the current security system namely, intrusion detection and prevention system, firewall, and internal vulnerability eradication program. The CIO's objective is to gauge if the company is spending too much or too little in security based on the return on investment measure. She knows that key to her task is determining the effectiveness of these components in securing the network against security incidents.

For a span of 12 months, she and her team recorded network security incidents and classified them into three types: type A for root compromises (e.g. hacker gaining root access to a user account); type B for malicious code infections (e.g. worms and viruses); and type C for improper use (i.e. leaks and potential embarrassment to the organization). The total damage for all recorded incidents for the 12-month period was estimated to be $5,000, $6,000, and $4,000 for type A, B, and C incidents respectively (see Exhibit 3). These damages include resources used to repair damaged information, isolate and remove any errant data in the network, bringing back the system to previous state, and any lost productivity. She realizes that these recorded incidents are only the tip of the iceberg. If the security components are as good as many thinks, then there are plenty more misses and near-misses that failed to be recorded.

After close collaboration with vendors and IT professionals, her team was able to estimate how effective or more accurately, how ineffective the security components are in preventing incidents. The team noticed that firewalls and vulnerability eradication do not provide security against type C incidents while the intrusion detection and prevention system is ineffective 10 per cent of the time. The team decided to use the notion of bypass rate of each types of incident which describes percentage of incidents that were able to bypass a particular security component, and thus gets recorded. Since there is no way to ascertain the interaction of the three security components, the net bypass rate of each incident type was assumed to be the product of the rates for the three security components (see Exhibit 3).

The CIO then wants to get a picture of the actual risk the various incidents create - the damage if there were no security in place. This was obtained by dividing the total damage for each type of incidents by the corresponding net bypass rate. As an example, consider type A incidents with a total damage of $5,000, the result of the net 0.225% which was able to bypass the three security components. The incident risk for type A incident is $5,000/0.225% = $2,222,222. In the absence of all three security components, this is how much damage the company would have suffered. The CIO was startled that even though the damage brought by the incidents does not significantly vary among the types of incidents, their risk definitely does (see Exhibit 3). Now that the CIO is aware of how much risks all the incidents create - a perplexing sum of $6,262,222 for all types of incidents - she wants to know each component's contribution in reducing these risks. Basically, she wants to know the residual risks - the risks if only one component is in place. Computationally, the residual risk is simply the sum of the products of incident risk and the corresponding bypass rate. For Intrusion detection and prevention, the residual risk is $2,222,222*10%+$4,000,000*10%+$40,000*10%=$626,222. Thus, this particular component has a benefit of reducing total risk by $6,262,222 - $626,222 = $5,636,000. Its net benefit (i.e. benefit minus implementation cost) is $5,636,000-$300,000 = $5,336,000 (see Exhibit 4).

The CIO now has useful information on hand: the net benefit for each of the security components. Together with the cost of implementation, she can perform cost-benefit analysis. For intrusion detection and prevention, the risk-based ROI is its net benefit divided by its implementation cost: $5,336,000/$300,000 = 18.

Exhibit 3 shows the RROI for individual security components and for the three all in place. With the company policy of investing only if the return is at least 10%, the CIO is convinced that the system made up of the three components is a worthy investment. Furthermore, she is confident that the company is neither investing too much or too little in security. However, she also recognizes limitations of the process, particularly the possible unintended effects of aggregating incidents to types. There obviously is heterogeneity even among incidents that belong to the same type, and representing all of them by a single damage estimate may bring inaccuracies. However, she recognizes the advantage of having a numerical RROI to aid in investment decisions in risk management.

Exhibit 3. Bypass Rates and Incident Risks.

	Types of incidents		
	A	B	C
Total damage	$5,000	$6,000	$4,000
Bypass rates:			
Intrusion detection & prevention	10%	10%	10%
Firewall	15%	15%	100%
Vulnerability eradication	15%	10%	100%
Net bypass rate	0.225%	0.15%	10%
Incident risk	$2,222,222	$4,000,000	$40,000

Exhibit 4. Residual Risks, Implementation Costs, Net Benefits, and RROI.

	Implementation cost	Net benefit	RROI
No security component	$0	($6,262,222)	--
Intrusion detection & prevention	$300,000	$5,336,000	18
Firewall	$75,000	$5,213,889	70
Vulnerability eradication	$200,000	$5,288,889	26
Entire security system	$575,000	$5,672,222	10

References

Bryson, Kweku-Muata and William E. Sullivan "Designing effective incentive-oriented outsourcing contracts for ERP systems," *Proceedings of the 35th Hawaii International Conference on System Sciences* (2002) pp. 2760 – 2769.

Clements, Paul C. and Ray C. Williams *Final Report From the Workshop "Managing Software Risk At NASA", 21-22 October 2003* CMU-Software Engineering Institute (November, 2003).

Fudenberg, Drew and Jean Tirole *Game Theory* The MIT Press, (2002) pp. 253-257.

Haimes, Yacov, *Risk modeling, assessment, and management* J. Wiley & Sons (1998).

Heimann, C. F. Larry Acceptable *Risks, Politics, Policy, and Risky Technologies* The University of Michigan Press (2000).

High Dependability Computing Project (HDCP) *HDCP Research Activities* http://hdcp.org/Research/index.html, (cited May 2004).

Kaplan, Stan and B.John Garrick, “On the quantitative definition of risk,” *Risk Analysis* Vol. 1, No. 1, (1981), pp. 11-27.

Schmitz, Edward, C. Ariel Pinto, Ashish Arora, and Rahul Telang, “Iterative Incentive Scheme for Outsourced IA,” *Proceedings of the 2004 IEEE Workshop on Information Assurance* (June, 2004).

Society for Risk Analysis (SRA), *Risk Analysis Glossary* http://www.sra.org/resources_glossary_p-r.php, (cited May 2004).

Varian, Hal R., *Microeconomic Analysis* 3rd edition W.W. Norton & Company (1992) pp. 440-445.

About the Authors

C. Ariel Pinto received his Ph.D. from the University of Virginia and his M.S. and B.S. from the University of the Philippines. He is an Assistant Professor of Engineering Management and Systems Engineering at Old Dominion University. His research focuses on risk management in engineered systems.

Ashish Arora received his Ph.D. from Stanford University. He is an Associate Professor of Economics and Public Policy at Carnegie Mellon University and serves as the co-director of the Software Industry Center, and was a founding co-director of the Sustainable Computing Consortium. His research focuses on the economics of technological change, the management of technology, intellectual property rights, and technology licensing.

Dennis Hall received his B.A. from the University of California at Berkeley. He helped found the Software Tools Users Group at Berkeley. He was awarded the Usenix Lifetime Achievement Award in 1996 for his pioneering work and outstanding contributions to the open-source software community.

Edward Schmitz received his Ph.D. from the University of Maryland in Economics (Policy Sciences), his M.S. from Carnegie Mellon University in Public Policy & Management, and his B.S. from Rensselaer Polytechnic Institute. He is currently serving as Special Assistant for Performance Management to the Director of the Navy Marine Corps Intranet Program.