

Improving Information Security Risk Management

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Anand Singh

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

David Lilja, Advisor

December, 2009

© Anand Singh 2009

Acknowledgements

As I reach the end of the journey towards my doctorate, I am filled with gratefulness towards so many whose direction, blessings, guidance and mentorship helped me accomplish this goal. I would like to take this opportunity to thank from the bottom of my heart everyone who helped me in pursuing this research.

I would like to thank my grandfather Shri Ram Roop Singh. He was a freedom fighter who did his part in making India independent. His lifelong emphasis on education created a generational shift in our family towards the better. Baba, your life story, words and actions taught me the virtue of hard work, inspired me to complete this Ph.D. and outfitted me with tools to deal with almost any circumstance. Thank you!

I am deeply indebted to my advisor, Prof. David Lilja. His vast technical knowledge was instrumental in seeding the idea as well as in bringing key ideas to fruition. His infinite patience helped me stick to this goal as my life circumstances changed and I had to extend the duration of this research. Most importantly, I have learned from him the leadership principles of integrity, clear communication and empathy.

I would like to thank the other members of my Ph.D. committee: Professors Abhishek Chandra, Antonia Zhai, Wei Chung Hsu and Gautam Ray. They provided invaluable feedback and advice for my work and this dissertation. I would also like to thank Professor George Karypis for guiding the initial part of my research.

I would like to thank various people who have touched my life over the last few years and helped me directly and indirectly with my thesis: Karl Baltes who has been my mentor since 2005 and every success I have had since then has been impacted by his advice and guidance; Bryan Koemptgen who showed me by example that most success is a result of nothing but hard work; Kurt Lieber for turning me into a true security

professional and for supporting my sabbatical so that I can work on my research and Anne Murphy for her outstanding advice on personal development. Nick, Jay, Adam: our lively discussions on risk management helped me understand the gaps in risk management methodologies addressing which is the crux of this thesis.

Good friends are a true blessing. I am fortunate enough to have a circle of friends who provided the much needed respite while I pursued this lofty goal. Dev, Meenakshi, Anu, Guru, Kalyani, Baiju, Sonia, Shail, Kishore, Muskan, Ravi, Shaili, Shubha, Venkat, Nisha and Anurag: thank you for being such a close part of our lives.

Completing a Ph.D. is a grueling proposition and impossible but for the support and encouragement of family. I would like to offer my heartfelt thanks to my parents Dr. Tara Devi Singh (she got her Ph.D. with three kids and everyday that I thought I couldn't do this, all I had to do was to look up to her to know that it is possible) and Dr. Ashok Kumar Singh, my brother Ajit and sister-in-law Priya, my sister Abhilasha and brother-in-law Vikas, my nieces Anika and Shreya and my nephew Amrit. Heartfelt thanks are also due to my in-laws Surendra Pratap Singh and Saroj Lata Singh who supported us in the time of our maximum need without which, this Ph.D. would not have been possible.

And finally, but most importantly, I would like to thank my wife Soma and my daughter Arushi. They supported me at times when I couldn't keep myself going and tolerated me when I spent 20 hour days engrossed in my research while abdicating my responsibilities as a father and a husband. The thought of seeing the pride on your faces as I receive my degree kept me going. Thank you for always being there for me, for your unrequited support and for being an inexhaustible source of happiness and affection. My son Arnav, who arrived in this world while I was writing this thesis, deserves special thanks as well: it seems as if he knew that Papa needed to focus on his thesis and so he did his best to be a calm, quiet baby!

Dedication

To my grandfather Shri Ram Roop Singh who instilled in me the values of the past and to my children Arushi and Arnav in whose eyes I see the brightness of the future.

Abstract

Optimizing risk to information to protect the enterprise as well as to satisfy government and industry mandates is a core function of most information security departments. Risk management is the discipline that is focused on assessing, mitigating, monitoring and optimizing risks to information. Risk assessments and analyses are critical sub-processes within risk management and are used to generate data that drive organizational decisions to accomplish this objective. However, despite this need, current approaches lack granular guidance on some key steps and have focused on qualitative data rather than quantitative data which reduces the value of the results for the decision makers. Through our research, we have identified the gaps in existing risk management methodologies. We have developed statistical design of experiments and requirements engineering based approaches to address these gaps. In addition, our quantitative models lead to a better alignment with business objectives by providing data to address the economics of making security decisions. Towards these ends, the work proposed here comprises of the following key components:

- (a) Improving risk assessment methodology through statistical models for control subsetting, configuration determination and judging the impact of security enhancements.

- (b) Developing approaches for dynamic configuration adjustment in response to changing security posture of an enterprise.
- (c) Managing the information risk introduced by vendors of an enterprise
- (d) Using requirements engineering to develop criteria and methodology for governance, risk management and compliance (GRC) which are used to drive risk considerations across the enterprise.

Our research makes extensive use of statistical models; specifically, we are using Plackett-Burman statistical design of experiments technique for prioritizing security controls. Once prioritized controls have been determined, we propose the usage of control sensors to dynamically recommend security configuration adjustment. We also intend to use requirements engineering to develop process frameworks for managing security risks introduced by the vendors of an enterprise as well as for GRC management.

Table of Contents

Abstract.....	iv
List of Tables.....	ix
List of Figures	x
1. Introduction	1
1.1 Background.....	3
1.2 Desirable Characteristics in a Risk Management Methodology	5
1.3 Improving Risk Assessment Methodology and Defining Quantitative Measures	6
1.4 Managing Risks Associated with Vendors.....	9
1.5 Governance, Risk Management and Compliance.....	9
1.6 Contributions of this Dissertation.....	10
1.7 Research Resources	11
2. Related Work	13
2.1 Enterprise Information Risk Management	13
2.2 COBIT	19
2.3 Security Parameter Determination.....	20
2.4 Screening Designs.....	21
2.5 Requirements Engineering for Proactive Risk Management.....	21
3. Improving Risk Assessment Process.....	24
3.1 Background.....	24
3.2 Plackett-Burman Design of Experiments	27
3.3 Threat Classification.....	30

3.4	Criteria for Determining Cost of a Security Incident	31
3.5	Determining PB High and PB Low Values	32
3.6	Experiment Setup	32
3.7	Determining Critical Controls	34
3.8	Determining Configuration of Controls	38
3.9	Determining Impact of Security Enhancements	39
3.10	Summary.....	42
4. Dynamic Security Configuration Management.....		44
4.1	STARTS Architecture	45
4.2	Experiments and Results	50
4.3	Summary.....	55
5. Managing Security Risks Due to Vendors.....		57
5.1	Security Risks Introduced by Vendors	57
5.2	VIAP Framework Overview.....	60
5.3	Vendor Security Governance.....	62
5.4	Due Diligence on Vendor Contracts.....	66
5.5	Measuring Vendor Compliance.....	67
5.6	Implementation Results	72
5.7	Summary.....	74
6. Governance, Risk Management and Compliance Platform Selection.....		76
6.1	General Considerations.....	78
6.2	Functional Requirements.....	79
6.3	Non-Functional Requirements.....	83
6.4	Selection Process Walk-Through	85
6.5	Summary.....	86
7. Conclusion.....		90

8. Future Work.....	94
8.1 Dynamic Security Configuration Management.....	94
8.2 Vendor Risk Management.....	95
8.3 Quantitative Measures for Security Decisions	95
References	97
Appendix A -- Permissions to Reprint	105
Vita	106
Listing of Research Output.....	107

List of Tables

Table 3.2.1	Design Matrix for N=8 with Foldover.....	29
Table 3.3.1	Threat Classification used in Experiments Conducted.....	31
Table 3.4.1	Criteria for Determining Cost of a Security Incident	31
Table 3.5.1	PB Values for Experiment Security Controls	33
Table 3.6.1	PB Design Results for all Security Controls Sorted by Sum of Ranks...	36
Table 3.8.1	PB Design Results Post Implementation of Security Code Review Recommendations	40
Table 5.5.1.1	Criteria for Risk Tiering	67

List of Figures

Figure 2.1.1.1	NIST Risk Assessment Process.....	14
Figure 2.1.1.2	NIST Risk Mitigation Process.....	15
Figure 2.1.2.1	ISO 27005 Information Security Risk Management Process.....	16
Figure 2.8.1	The Integration of Risk Management with SDLC as proposed by NIST	22
Figure 4.1.3.1	The STARTS Architecture	47
Figure 4.3.1	Detailed Results Comparing before and after STARTS Implementation.....	53
Figure 5.3.1	Sample Vendor Information Security Policy	63
Figure 5.3.2	Sample Vendor Information Security Standard	64
Figure 5.3.2	Sample Vendor Information Security Standard (Contd.).....	65
Figure 5.5.2.1	Flowchart depicting steps in vendor compliance audit	70
Figure 5.6.1	Comparison of Resource Hours before and after VIAP Implementation.....	72
Figure 5.6.2	Comparison of Incidents Before and After VIAP Implementation....	73
Figure 6.1.1	Evaluating general considerations	78
Figure 6.2.1	Governance Requirements.....	81
Figure 6.2.2	Risk Management Requirements.....	81
Figure 6.2.3	Compliance Requirements.....	82
Figure 6.4.1	Requirements Solicitation Questionnaire	88
Figure 6.4.2	Criteria Weight Determination	88
Figure 6.4.3	Decision Table.....	89

Chapter 1

Introduction

Peter Drucker once said “the diffusion of technology and commoditization of information transforms it into a resource equal in importance to the traditionally important resources of land, labor and capital” [14]. The exponential growth and availability of information after the Internet boom of 1990’s goes to show the accuracy of his foresight. In today’s world, the fortunes of most organizations are tied with the information they possess and the sophistication with which they are able to manage it.

As a consequence of information becoming central to organizational strategies, risk management of information is increasingly being rolled into the organizational risk profile. Following factors are also increasing the focus on managing risks around information:

- An explosion in eCommerce [2]
- Outsourcing of IT Assets [6]
- Offshoring of IT Assets [4]
- Mobile Computing [28]

Information risk management is the activity directed towards assessing, mitigating (to an acceptable level) and monitoring of risks associated with information. The principle goal of an organization’s risk management process should be to protect the *organization and its ability to perform their mission*, not just its IT assets [70].

In the last few years, a number of risk management methodologies have been developed, both in the academic as well as commercial sectors. These have shed insights into the problems that still need to be solved before we can have risk assessments and analytics that are useful and can be used to satisfy the economics of decision support systems in information security. As a result, this area has started seeing increasing research with request for proposal's being routinely sought by government and technology companies to solidify the gaps that exist. As an example, HP has recently solicited proposals in security analytics to *"bring economic, mathematical, and cognitive modeling to bear on systems understanding; improving the framework for making security decisions by applying modeling techniques"*.

Most of these risk management methodologies, while providing a structured and systematic process for risk management, either lack specific guidance on which risk assessment methods to use or provide for a weak approach. This does not satisfy the rigorous data needs of business leaders as well as audit needs of compliance auditors. This was clearly identified as a significant issue in the recent RSA report based on discussions with top risk management leaders in Global 1000 companies [61]: *"Risk should be managed to an acceptable level, based on the enterprise's risk appetite with decision-making guided by a risk assessment model. A structured, consistent and repeatable process for making the risk/reward calculation helps to ensure that it is done consistently across the organization"*.

However, despite this need, it remains un-addressed in current state of the art:

- A risk assessment typically exercise involves the following steps:
 - (a) Identification of controls to be tested.
 - (b) Testing of these controls for their efficacy.

(c) Analysis of test results.

(d) Recommendations for security enhancements based on analysis.

While providing a clear process framework, most risk assessment methodologies lack clear guidance on how these steps can be accomplished in a structured fashion and with consistency.

- Another key question that remains unanswered in the current research is how an enterprise can adapt to the changing risk environment around it. This is an important question because a static security configuration loses its utility in a very short timeframe given the dynamic nature of threats.
- In today's environment, enterprises are significantly dependent on their vendors for non-core function. Thus vendors tend to have access to information sensitive to the enterprise. However, no formal means exist to evaluate and manage this risk to the enterprise.
- Finally, risk considerations are driven through the enterprise using governance, risk management and compliance (GRC). However, there is a dearth of formal means to determine how models for GRC should be selected.

1.1 Background

We start by introducing concepts that are used throughout this thesis.

Risk Management

The dictionary definition of risk management is that it is the activity directed towards assessing, mitigating (to an acceptable level) and monitoring of risk. Another definition

[20] says, “risk management, in general, is a process aimed at an efficient balance between realizing opportunities for gains and minimizing vulnerabilities and losses”. Given that risk/reward equation associated with information risk is the need of the hour [61], the second definition is what we would strive for in this research.

Risk Assessment

Risk assessment is the determination of quantitative or qualitative value of risk related to information.

Security Controls

Activities or technology solutions that address risk (or mitigate it to an acceptable level).

Governance

The set of responsibilities and practices exercised by the enterprise board of directors and executive management with the goal of providing strategic direction, ensuring that the objectives are achieved, ascertaining that risks are managed appropriately and verifying that the enterprise’s resources are used responsibly.

Compliance

Compliance is either a state of being in accordance with established guidelines, specifications, or legislation (e.g. GLBA, HIPAA, SOX, PCI etc.) or the process of becoming so.

Policy

High level statement of executive management's intent or direction.

Standards

Standards are the metrics, allowable boundaries or the process used to determine whether processes meet policy requirements.

Procedures

A detailed description of the steps necessary to perform specific operations to achieve conformant with applicable standards.

Guidelines

A suggested action or recommendation related to an area of information security policy that is intended to supplement a procedure. Unlike standards, implementation of guidelines may be at the discretion of the reader.

1.2 Desirable Characteristics in a Risk Management Methodology

A comprehensive definition of the characteristics desired from a risk management system in one place is missing from current literature on this topic. We propose the following criteria based on our research (these criteria are articulated in our paper [64]):

1. It must manage risks to an *acceptable* level based on enterprise's risk appetite [61].

2. It must provide *decision-support* [61]. Security investments are expensive and risk is one criterion that is used to address the economics around it.
3. It must be a *continual* process [35]. Risk management is not conducted at a point in time; it should be considered throughout the lifecycle of systems development.
4. It must be *aligned* with an organization's business objectives [70].

As the amount as well as complexity of information resources within organizations is increasing at an exponential rate, we also consider the following characteristics as desirable traits:

5. It must be *adaptive*. Since an organization's risk profile, threats and vulnerabilities change frequently, it is important that risk management should be adaptive to these changes.
6. It must be *scalable* to accommodate for this increasing complexity while not impacting the window desired to conduct the assessment activities.
7. It must ensure *compliance* with government and industry mandates.
8. It must produce *consistent* results irrespective of who conducts the responsibilities associated with risk management.

1.3 Improving risk assessment methodology and defining quantitative measures

Risk assessment is a critical component of risk management process. Existing risk assessment and analytics methodologies establish good process frameworks but have several gaps as far as implementation specifics are concerned. Specifically, they tend to focus on qualitative data thus reducing its value to the decision makers [5]. In addition, they do not provide for the means to meet the desired objectives as defined in Section 1.2. Our research is focused on identifying these gaps and addressing them so that the desired characteristics of a risk management model can be met. Our proposed enhancements

make the risk management framework scalable, proactive and quantitative thus improving enterprise risk management and meeting business leaders' need for smart security decision-making and to balance the risk/reward equation.

Following sections describe the gaps we have identified through an analysis of the prevalent risk management methodologies as well as how it integrates in the overall risk management model:

Existing risk assessment approaches suffer from the following issues:

1.3.1 Identification of Critical Controls

Testing every control for the purposes of risk assessment is a very expensive and time consuming proposition. Therefore risk assessors will often prioritize controls on the basis of criticality and focus on the most critical controls of an enterprise when conducting risk assessments. However, no formal models exist as far as how to conduct this prioritization; oftentimes it is conducted informally with the enterprise identifying the important controls or the risk assessor using his/her own judgment.

1.3.2 Configuration of Critical Controls

Different enterprises will face different threats. For e.g. a web commerce site is likely to be more impacted by a denial of service attack as compared to a research institution where theft of intellectual property would be primary concern. Clearly, the security controls of an enterprise need to be configured based on the threats faced by it as well as the cost of fruition of those threats .While some preliminary research has taken place in this area, the models are niche models which cannot be scaled to generally usage across the board. Also, these models are not quantitative in nature and do not take into account the risk appetite as well as the costs of threats coming into fruition. So they don't address

the data needs of the key decision makers of the enterprise. A structured, consistent and quantitative approach for determining security configuration based on cost to the enterprise and the threats faced by it is the need of the hour.

1.3.3 Impact of Security Enhancements

Based on the risk assessment analytics, a risk assessor provides recommendation on how controls need to be adjusted or whether new controls need to be added. However, decision makers want to measure the impact of these security enhancements. For e.g. increasing the strictness of configuration of a control might mean that the end user sees increased response times; for a decision maker, it is critical to understand whether increasing that strictness and the inconvenience caused to the end user as a result is worth it or not in terms of prevention of security threats. While this area has been researched in other disciplines such as microprocessor simulation [78], it remains unaddressed within the domain of information security.

1.3.4 Dynamically adapting to security threats

The risk profile of an organization changes on a very dynamic basis because new threats come into existence on an almost continuous basis. Thus any approaches to deal with the threats have to be dynamic as well. This issue has not been dealt with in existing research, either methodologically or architecturally.

We are proposing the design and implementation of architectural concepts called *control sensors* to handle this aspect of risk.

1.4 Managing Risks Associated with Vendors

Current risk management approaches often lack specific guidance on how to be proactive about managing risks to the enterprise. Most IT organizations are exposed to significant risk from their vendor providers because of the vendor's risk to enterprise's data and infrastructure. As an example, 401k vendor to the enterprise will have access to confidential employee data. Another example is that of an offshore development firm which will have network connectivity to enterprise infrastructure as well as access to critical server resources.

Despite this significant risk introduced by the vendors, these areas have not been accounted for in the organizational risk profile. In our research, we propose a process framework using requirement-engineering approach to define the security requirements upfront and use them to build SLA's so that the proprietary data of the enterprise can be protected and the risk of these relationships to the enterprise reduced.

1.5 Governance, Risk Management and Compliance

Managing governance, risk and compliance across the enterprise is a challenging proposition. To address this issue, a class of solutions known as GRC (governance, risk management and compliance) have sprung up. In addition, many enterprises will end up developing and implementing custom solutions that they build in-house. However, coming up with the right requirements for the appropriate GRC solution is a complicated undertaking. Every enterprise will have unique needs that will drive what the GRC solution should look like. These needs might be the compliance regiment that the enterprise has to adhere to (e.g. PCI), the threats faced by it, its risk appetite or a multitude of other considerations. No formal models exist that help the definition of an enterprise's GRC needs. We address this gap by providing selection criteria and methodology for GRC platform selection in Chapter 6.

1.6 Contributions of this Dissertation

Our foundational research identifies the gaps in existing risk management approaches and clearly demonstrates the need for methodological improvement in those to further appropriate information security decision making, resource optimization and risk management. It also establishes the criteria that can be used to evaluate these approaches. We provide specific recommendations on how to improve risk assessment methodology through the following means: (a) a statistical model to identify critical controls of an enterprise based on the threats faced by it; (b) a structured approach for determining control settings at a macro level; and (c) an approach for determining the effect of changing control settings. By way of illustrating and validating (b), our research determines the critical controls for a product data management (PDM) production system and compares it with the critical controls identified through a risk assessment audit conducted by trained security experts on the same system.

One key contribution of our research is a new requirements engineering approach that will enable the binding of organizational security policies and standards to the governance, risk management and compliance requirements. This contribution changes the landscape of risk management from *reactive* to *proactive* thereby addressing the second opportunity identified in Section 1.2.

Our research combines multiple techniques (statistical design of experiments, past research on the cost of a security breach, Center for Internet Security benchmarks, vendor benchmarks and web security threat categorization) to improve the risk assessment methodology. This furthers the *consistency* and *scalability* objective identified in Section 1.2. This is a very significant contribution because it addresses gaps in existing methodologies.

An innovative aspect of this approach is the use of control sensors to automate the task of risk assessment and enable dynamic adjustment in response to threats impacting the enterprise. In addition, the use of control sensors will generate data and patterns over time that can be processed to produce detailed, quantitative data that is needed to make complex security decisions thus addressing issues 4 and 5 identified in Section 1.2. The use of control sensors is an innovative contribution for another reason as well – they can be used to extract the data needed to demonstrate compliance with appropriate security standards and organizational goals. This would be especially valuable for large organizations that are spread across extensive geographies with different rules within local jurisdictions. This furthers the *compliance* objective.

1.7 Research Resources

1.7.1 Resources Provided by the University of Minnesota

The areas of security analytics and risk management require extensive interdisciplinary collaboration. We have leveraged University of Minnesota’s extensive research contributions from related departments and their services towards this end. We had partnerships with the following areas that strengthened our ability to deliver on our research.

- School of Statistics and their Statistical Consulting Services
- Carlson School of Management
- Other faculty in the Department of Electrical and Computer Engineering and the Department of Computer Science and Engineering
- Computing resources from the departments and Minnesota Supercomputing Institute

In addition to University of Minnesota’s research facilities, we also took advantage of extensive research conducted in this area at CERIAS (Center for Education and Research

in Information Assurance and Security) labs at Purdue University, which is a premier research institution in the area of information security. The author spent several days onsite at the lab in an exchange of ideas on my research topic. The input provided by CERIAS researchers was instrumental in solving some issues encountered during the conduct of the research.

1.7.2 Shared Assessments Group

Shared Assessment Group (SAG) [24] is an industry body responsible for developing uniform standard for managing information risk introduced by vendors to an organization. Since one facet of our research (policies for managing vendor risk) overlaps with SAG's interests, they have agreed to provide consulting services as well as facilitate interviews and collaboration with SAG's members.

Chapter 2

Related Work

As described in Chapter 1, Risk Management is risk management, is a process aimed at an efficient balance between realizing opportunities for gains and minimizing vulnerabilities and losses. The remainder of this chapter describes currently prevalent risk management models, the gaps therein and the current state of the art in addressing those gaps.

2.1 Enterprise Information Risk Management

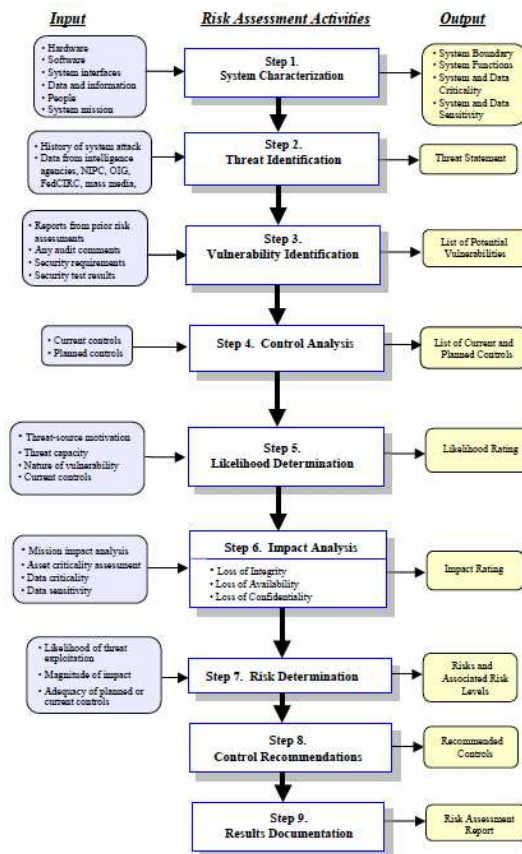
A considerable body of work is available on the process aspects of risk management from industry bodies and standards organizations. While they have covered the issue of defining the process around risk management, the granular guidance on the issues identified in Chapter 1 is missing. These models have typically focused extensively on the issue of defining a process around risk management. While excellent from a process perspective, these models have either not defined rigorous means for *how* to accomplish some key steps (the topic of this research). Following is a description of key risk management models prevalent in various sectors.

2.1.1 Risk Management Guide for Information Technology Systems

The National Institute of Standards and Technology (NIST) provides a foundation for the development of a risk management program and is particularly prevalent in the government sector [49]. Their *Risk Management Guide for Information Technology Systems* [70] is considered the seminal work in this area. It is credited with establishing

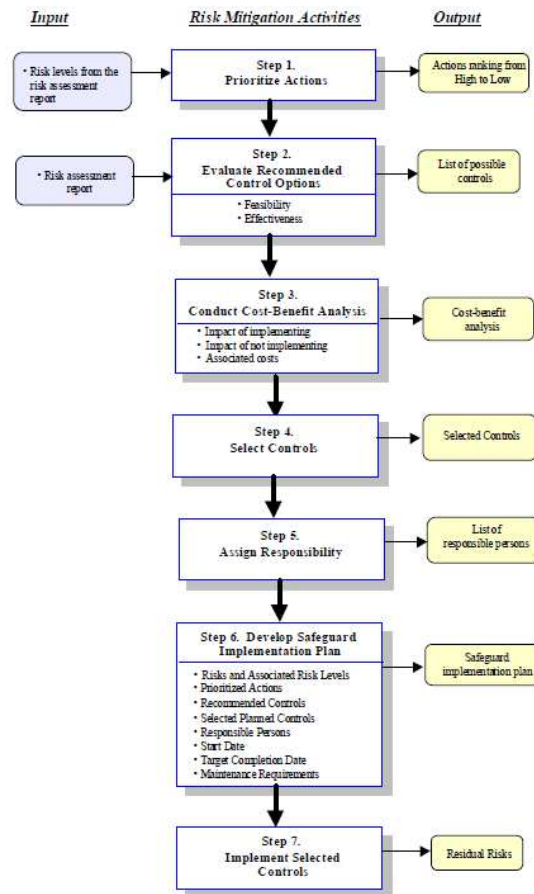
widely accepted definitions of key risk terms. This document is primarily a guidance document, though, geared towards informing risk managers about the key concepts needed for risk management at a high level. This document identifies “Risk Assessment” and “Risk Mitigation” as two crucial components of Risk Management. Fig. 2.1.1.1 summarizes their “Risk Assessment” process and Fig. 2.1.1.2 summarizes their “Risk Mitigation” process. In Risk Assessment process, Control Analysis is Step 4. However, it does not provide more granular guidance on how “Control Analysis” should be conducted. Similarly, the risk mitigation process identifies “Control Selection” as one of the key steps. However, no further guidance is provided on how controls should be selected for testing purposes.

Fig. 2.1.1.1 NIST Risk Assessment Process ¹



¹ Fig. 2.1.1.1 reprinted from [70] courtesy of the National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce. Not copyrightable in the United States.

Fig. 2.1.1.2 NIST Risk Mitigation Process ²



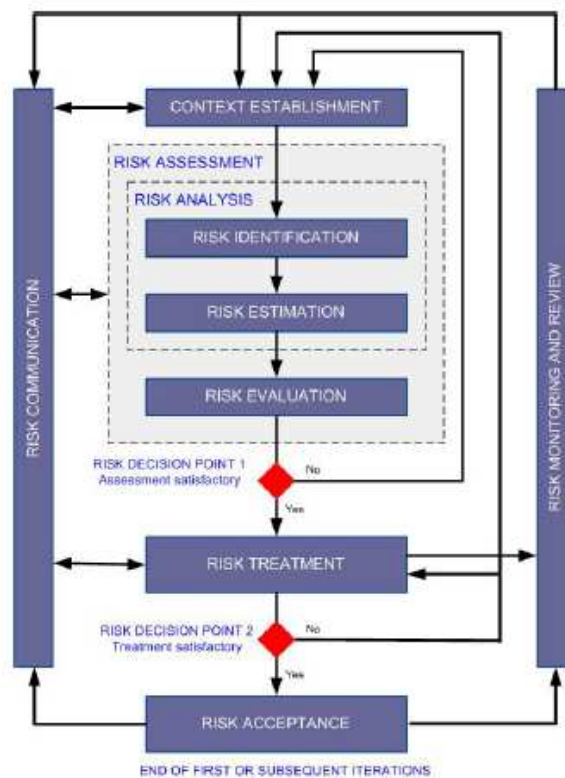
2.1.2 ISO 27005

ISO 27005 provides guidelines for Information Security Risk Management in an organization, supporting in particular the requirements of an Information Security Management System (ISMS) [35]. However, this international standard does not provide any specific methodology for information security risk management. It provides a process framework and leaves it up to the organization to define their approach to risk management, depending on niceties like ISMS or the context of risk management.

² Fig. 2.1.1.2 reprinted from [70] courtesy of the National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce. Not copyrightable in the United States.

Figure 2.1.2.1 shows the summary ISO 27005 process framework. It is clear from this diagram that in this approach, information security risk management process can be iterative for risk assessment and/or risk treatment activities. An iterative approach to conducting risk assessment can increase depth and detail of the assessment at each iteration. The iterative approach provides a good balance between minimizing the time and effort spent in identifying controls, while still ensuring that high risks are appropriately assessed. Like NIST, ISO 27005 fails to provide granular guidance on key steps of critical control identification and configuration.

Fig. 2.1.2.1 ISO 27005 Information Security Risk Management Process ³



³ Fig. 2.1.2.1 reproduced from ISO/IEC 27005 [35] with permission of the American National Standards Institute (ANSI) on behalf of the International Organization for Standardization (ISO). No part of this material may be copied or reproduced in any form, electronic retrieval system or otherwise or made available on the Internet, a public network, by satellite or otherwise without the prior written consent of the ANSI. Copies of this standard may be purchased from the ANSI, 25 West 43rd Street, New York, NY 10036, (212) 642-4900, <http://webstore.ansi.org>

2.1.3 OCTAVE

OCTAVE [50] from the CERT program at Carnegie Mellon University goes much further than NIST and ISO 27005 in terms of providing prescriptive guidance and diving into the details of risk management and not just providing a process framework. It is intended to help an organization develop qualitative risk criteria that describe its operational risk tolerances, identify assets that are important to its mission, identify vulnerabilities and threats to those assets and determine consequences to the organization if the threats are realized. Risk management is covered in OCTAVE's Strategy and Plan Development Phase [51, 52, 53]. This phase includes the following outputs:

- Risks to Critical Assets
- Risk Measures
- Protection Strategy
- Risk Mitigation Plans

OCTAVE's focus is primarily on qualitative risk determination [55]. As is evident from [61], however, the landscape has shifted substantially towards the need for more quantitative information with which business leaders are familiar for decision-making purposes.

2.1.4 IRAM

IRAM [81] is a proprietary methodology available from the Information Security Forum that uses business impact analysis approaches for risk analysis. It is geared towards assessing the business impact of potential security breaches, assessing threats and vulnerabilities, determining information risks, identifying and analyzing control

requirements and generating an action plan to address the identified control requirements. IRAM risk assessment process has three fundamental steps:

- **Business Impact Assessment:** This sub-process is focused on determining the impact to the business if the system were to become un-available.
- **Threat and Vulnerability Assessment:** This sub-process starts with determination of the system profile and laying out a detailed plan for the assessment. Once the plan has been laid out, an assessment for threats and vulnerabilities is conducted.
- **Control Selection:** The steps in “Control Selection” sub-process are judgment based with qualitative interviews with key stakeholders being used as the primary mechanism for risk assessment. As has been demonstrated in [59], judgments are not accurate means to assess risks and frequently, stakeholders might have conflicting opinions about the criticality of an asset.

2.1.5 Problems with these Risk Management Models

As is evident from the risk management methodologies discussed in previous sections, Risk Assessment is a crucial part of risk management. Within the risk assessment process of all the models described above, critical control identification and configuration are critical needs. In addition, OCTAVE and IRAM are qualitative in nature and ISO 27005 and NIST operate at a process framework level and hence, don't address the quantitative needs of decision makers [70]. These issues expose the following problems in current risk management approaches:

- Risk assessment is a critical component of risk management exercise. Existing risk assessment approaches are long drawn out exercises that don't suffice for the times when there is a need to conduct risk assessments very rapidly (e.g. in situations where a threat assumes alarming proportions in a short timeframe). They do not provide for mechanisms to customize control sets used for risk assessment depending on the needs and risk appetite of the organization. A security prioritization model that takes

into account the criticality of security controls would greatly alleviate these issues. In addition, configuration of security controls is an un-addressed problem in the current state of the art.

- Another un-addressed area in current risk management approaches is how to manage changes in security configuration if the nature of threats changes (e.g. if security controls are sensing a change in threat such as a denial of service attack, how should the security control configuration change?).
- Oftentimes, there is a need to conduct risk assessments very rapidly (e.g. an impending threat such as a denial of service attack). The approaches presented above take a very significant amount of time and do not provide for means to customize control sets used for risk assessment depending on the immediate needs of the enterprise.
- Security controls are expensive and the funds available to address risks are limited. A means to measure the impact of implementing a security control within the enterprise would greatly further the economics of security decision-making. This quantitative aspect of security has not seen much research in current state of the art.

These areas have been explored in greater detail in future chapters. Chapter 3 addresses issues identified in (1), (2) and (3) whereas Chapter 4 addresses the issue identified in (3). Current state of the art in these areas is discussed in respective chapters as well.

2.2 COBIT

COBIT (Control Objectives for IT and Available Technology) [17] provides the authoritative and complete body of control sets that can be exercised to conduct risk assessments. It has 34 high level processes that cover 210 control objectives categorized

in four domains: Planning and Organization, Acquisition and Implementation, Delivery and Support, and Monitoring and Evaluation. We intend to leverage these processes and control objectives extensively in conducting the experiments necessary to test out our hypotheses.

2.3 Security Parameter Determination

Han et al [31] describe a fuzzy adaptive configuration determining system to detect and drop forged reports, which are trying to deceive the base station in sensor networks. This scheme applies to the narrow purpose of deception avoidance in sensor networks; it is not applicable in the scenarios where configuration parameters might interact with each other, as is the case with security controls. Singh and Lilja [65] have demonstrated that when the number of parameters under consideration is small, it might be feasible to conduct an ANOVA analysis to determine the value of configuration parameters. This research would be exploited in Chapter 3 to determine the configuration of critical controls once they have been determined through PB design.

A similar situation is encountered when computer architects have a need to simulate only a subset of benchmarks in a benchmark suite. A solution to this problem is proposed in [79] using a statistical design of experiments approach. It demonstrates a Plackett and Burman based approach to reduce the number of benchmark programs needed to assess a new processor design. Given that the desire to quantitatively understand the impact of specific controls on information risk is a fairly analogous situation, we have explored statistical modeling techniques as the means to identify a representative control set.

The area of dynamic security configuration determination has started seeing some research. Stephenson [69] proposes a colored petri net approach based upon modeling against global norms and quantification using statistical methods. However, the results from [69] did not demonstrate this method to be highly effective and generally

applicable. The area of proactive handling of security when developing applications is addressed through context aware security architecture in [19]. In their proposal, the authors argue that it is possible to collect the contextual information from resources through this context aware framework, which can then be used to drive the access control paradigm of risk management. However, this work is limited to access control paradigm only which is only one of a large number of disciplines in security.

2.4 Screening Designs

Since one of the key risk assessment objectives is to identify critical controls, it is pertinent to understand the current state of the art in screening designs. Trocine and Malone [72] compare statistical screening methods and recommend that the choice of which method to use should be based on number of variables, the behavior of those variables, the accuracy desired and the cost of conducting experiments. We evaluated the designs we considered against these criteria. Specifically we chose PB designs as the best fit for our problem because they require linear number of experiments, are monotonic in nature and interact but to a small degree, exactly the type of problem best solved by these designs.

2.5 Requirements Engineering for Proactive Risk Management

NIST provides a high-level process framework on how security can be integrated with software development lifecycle (SDLC). Their approach is shown in Figure 2.71. However, this guidance is provided at a generic level only and does not address the specifics of how various steps can be accomplished. Meyer et.al. [42] describe a requirements engineering framework for proactive risk management. They define the set of concepts and relationships taking place in the IS security risk management (ISSRM) within a UML class diagram. They also define a modeling language for this purpose.

Fig. 2.8.1 Integration of Risk Management with SDLC as proposed by NIST ⁴

SDLC Phases	Phase Characteristics	Support from Risk Management Activities
Phase 1—Initiation	The need for an IT system is expressed and the purpose and scope of the IT system is documented	<ul style="list-style-type: none"> Identified risks are used to support the development of the system requirements, including security requirements, and a security concept of operations (strategy)
Phase 2—Development or Acquisition	The IT system is designed, purchased, programmed, developed, or otherwise constructed	<ul style="list-style-type: none"> The risks identified during this phase can be used to support the security analyses of the IT system that may lead to architecture and design trade-offs during system development
Phase 3—Implementation	The system security features should be configured, enabled, tested, and verified	<ul style="list-style-type: none"> The risk management process supports the assessment of the system implementation against its requirements and within its modeled operational environment. Decisions regarding risks identified must be made prior to system operation
Phase 4—Operation or Maintenance	The system performs its functions. Typically the system is being modified on an ongoing basis through the addition of hardware and software and by changes to organizational processes, policies, and procedures	<ul style="list-style-type: none"> Risk management activities are performed for periodic system reauthorization (or reaccreditation) or whenever major changes are made to an IT system in its operational, production environment (e.g., new system interfaces)
Phase 5—Disposal	This phase may involve the disposition of information, hardware, and software. Activities may include moving, archiving, discarding, or destroying information and sanitizing the hardware and software	<ul style="list-style-type: none"> Risk management activities are performed for system components that will be disposed of or replaced to ensure that the hardware and software are properly disposed of, that residual data is appropriately handled, and that system migration is conducted in a secure and systematic manner

Their research provides a formal way to tie security requirements within the UML development process. They further advance their research in [43] by defining the concept of misuse cases within UML to manage security risks. They use ISSRM modeling language to capture misuse cases. NIST also provides the model shown in Figure 2.8.1 for integrating risk management into software development lifecycle (SDLC) [70]. None of these papers discuss how governance, risk management and compliance requirements can be proactively managed within the enterprise. In addition, they don't tie the security

⁴ Figure 2.8.1 reprinted from [70] courtesy of the National Institute of Standards and Technology, Technology Administration, U.S. Department of Commerce. Not copyrightable in the United States.

requirements with organizational policies, standards or procedures or address the specific aspect of how the security risk introduced by the vendors to the enterprise (which can be very significant as exemplified by several recent security breaches caused by vendors to the enterprise such as BillPay and CareerBuilder).

Chapter 3

Improving Risk Assessment Process

As is clear from Chapter 2, Risk Assessment is a critical process within risk management. Typically it entails identification of controls to be tested, testing of those controls, analysis of the test results and recommendations on how these controls need to be configured based on the results of this testing [8]. However, the current state of the art does not provide for a formal model on how these three steps need to be accomplished. This is the problem that we explore and propose solutions for in this chapter.

3.1 Background

3.1.1 Identification of Critical Controls

Conducting risk assessment and audit of security infrastructure for a business process involves either testing the efficacy of *every* security control or prioritizing the controls and testing those deemed as critical [58]. Testing every security control is ideal but not practical in most circumstances because it is extremely time and resource intensive (testing every single control for a modestly complex business process can take several weeks of dedicated time from risk assessors). Given limited resources allocated towards this exercise in most organizations, it is also not possible in most situations. Also, it makes more sense to test critical controls for a larger number of business processes than to test all controls for just a few.

Therefore, in most cases, risk assessors use their subjective judgment to identify the critical controls. When the risk assessors belong to the enterprise, their subjective

judgments are likely to have a higher degree of accuracy since they are intimately familiar with the security architecture of the enterprise, past incidents as well as its risk exposure. However, more often than not, conduct of risk assessments is outsourced to outside experts because compliance regimes such as PCI, GLBA and HIPAA mandate that to maintain objectivity of the risk assessment exercise. Also, the ability to conduct risk assessments is a niche and expensive skill and hence, most organizations don't have that in-house. The judgment of an outside risk assessor on which security controls should be deemed critical is likely to have a lower degree of accuracy since they are not familiar with the specifics of the enterprise. Because of this lack of knowledge of specifics, an outside risk assessor is also less likely to understand the *interaction* between various controls in the security architecture thus further complicating their task of critical controls identification.

3.1.2 Configuration of Critical Controls

It is important to customize the controls configuration based on threats faced by the enterprise. This is because if all controls are configured for highest level of security, the resources costs for security, such as CPU time, become onerous and end-up impairing end-user satisfaction. If the controls are configured on the lower end, they might end up exposing the enterprise to threats. One common practice for controls configuration is to use the recommendations from the security solutions vendor or use security benchmarks from a industry standards body such as Center for Internet Security (CIS) [13]. However, such configurations are not likely to be *optimal* for most organizations because they are meant for *every* environment in which the solution would run and hence, are likely to have settings on the more stringent end thus leading to increased resource consumption.

Another approach some security architects use for determining appropriate controls configuration is to measure impact on costs while changing settings one at a time and keeping others constant. The central opportunity in this approach is that some controls

interact with each other because of overlapping functionality. Security controls might overlap with each other because of one or more of the following reasons:

- (a) Intentional (to ensure multi-layered security)
- (b) Different teams implementing different independent enterprise controls addressing the same risk
- (c) Loss of knowledge of existing controls over time leading to new controls being implemented
- (d) Enhancements in existing controls offering functionality of other controls. An example of such interrelationship is application-layer firewalls that have the ability to do some intrusion prevention system (IPS) like functions, such as enforcing RFC specifications on network traffic or doing signature analysis of incoming traffic. Given such interaction, it is not possible to determine settings for one control by keeping settings for other controls constant because it does not clarify what is the **impact** of the interaction.

3.1.3 Impact of Security Enhancements

One outcome of risk assessment exercise is recommendations on how to improve the security posture of the enterprise (e.g. introducing a new security control etc.). For decision-making purposes, it is important to understand in a measurable fashion, what impact these recommendations would have (e.g. if all recommendations cannot be implemented, it might make sense to implement only those that will have the most effect on improving the security posture). Current risk assessment methodologies do not provide specific guidance on how to measure the impact of making security enhancements.

Rest of this chapter is organized as follows: Section 3.2 describes our design of experiments methodology; Section 3.3 describes related work and how it will be used in our experiments; Sections 3.4, 3.5, 3.6 and 3.7 describe the experiments and their results; Section 8 has conclusions.

3.2 Plackett-Burman Design of Experiments

The parsimony principle says that in experiments, some of the factors are important while others are not [45]. Said another way, a few variables are responsible for most of the effect on the response while most variables contribute little [72]. As discussed in Section 3.1, more often than not, it is not possible to test every control at an enterprise. Therefore IT risk assessors frequently leverage the parsimony principle to examine key security controls at an enterprise for detailed testing and examination.

Another factor holds true about security controls is that they are monotonic in nature because if the control settings were gradually increased to a more stringent level, the security is likely to increase until it reaches a point at which any further increase does not have an effect on the security. Similarly, if the settings of the controls are diluted down, at a certain point, they will reach a level at which they provide no benefit to the enterprise.

Finally, while security controls do interact with each other, the degree of interaction is low. This interaction happens because of overlapping functionality among neighboring controls in the security architecture. As an example, Firewall and IDS/IPS are in close proximity to each other in a typical security architecture diagram and tend to have a small overlap in functionality [63]; however it is unlikely that a Firewall would end up overlapping in functionality with a File Integrity Monitoring (FIM) system which is several degrees away from it in a typical security architecture.

These factors make a screening design like Plackett & Burman (PB) designs ideal for determining the critical security controls [60]. These designs are very useful for economically determining the key variables most responsible for the end result. In fact, PB designs require logically minimum number of experiments to estimate the effect of each of the individual controls.

A PB design with X controls requires $X+1$ experiments where all X controls are varied simultaneously. Table 3.2.1 shows a PB design with $X=7$ controls (C_1 through C_7). Entries in rows indicate whether a high value (indicated by +1) or a low value (indicated by -1) has been chosen to configure that control. Since multiple controls are being varied simultaneously, it is important to choose the high and low values carefully so as not to have an outsize impact on the results. Typically, they are chosen to be slightly higher and slightly lower than what is considered normal. The fact that security controls tend to be monotonic helps because then, consideration of the two end values for the controls provides a good range of values to be considered without compromising the quality of the results.

The first row of the design matrix for other values of N ($N=8, 12, 16, \dots, 96, 100$), where N is the number of experiments is provided in [60]. The remaining rows can be constructed simply by circular right shift on the preceding row. As an example, row 2 in Table 3.2.1 is a circular right shift of row 1. This leads to an interesting situation where only $X+1$ experiments are required for X controls but the PB design for $X+1$ experiments will not exist if $X+1$ is not a multiple of 4. In this scenario, extra columns with dummy controls are added which have no effect on the overall results.

For experiments in this chapter, PB design with foldover has been used which is an improvement over the basic PB design described above. It introduces more rigor in the determination of the main controls and key interactions [44]. When using foldover, N additional rows are added to the design matrix. The gray area of Table 3.2.1 shows the foldover portion of the design matrix. The foldover rows are arrived at by reversing the

sign of the corresponding entry in the original matrix. As an example, row 9 is arrived at by reversing the signs in the entries for row 1.

Table 3.2.1. Design Matrix for N=8 with Foldover

	C_1	C_2	C_3	C_4	C_5	C_6	C_7	Cost
<i>Run</i> ₁	+1	+1	+1	-1	+1	-1	-1	21
<i>Run</i> ₂	-1	+1	+1	+1	-1	+1	-1	32
<i>Run</i> ₃	-1	-1	+1	+1	+1	-1	+1	19
<i>Run</i> ₄	+1	-1	-1	+1	+1	+1	-1	16
<i>Run</i> ₅	-1	+1	-1	-1	+1	+1	+1	44
<i>Run</i> ₆	+1	-1	+1	-1	-1	+1	+1	13
<i>Run</i> ₇	+1	+1	-1	+1	-1	-1	+1	11
<i>Run</i> ₈	-1	-1	-1	-1	-1	-1	-1	48
<i>Run</i> ₉	-1	-1	-1	1	-1	1	1	34
<i>Run</i> ₁₀	+1	-1	-1	-1	+1	-1	+1	48
<i>Run</i> ₁₁	+1	+1	-1	-1	-1	+1	-1	17
<i>Run</i> ₁₂	-1	+1	+1	-1	-1	-1	+1	8
<i>Run</i> ₁₃	+1	-1	+1	+1	-1	-1	-1	29
<i>Run</i> ₁₄	-1	+1	-1	+1	+1	-1	-1	33
<i>Run</i> ₁₅	-1	-1	+1	-1	+1	+1	-1	26
<i>Run</i> ₁₆	+1	+1	+1	+1	+1	+1	+1	3
<i>Control Effect</i>	-86	-64	-100	-48	+18	-32	-42	

After performing the experiments, the effect of each individual control is computed by multiplying the PB value with the result and then, summing those across all configurations. For e.g. the effect of C_1 can be computed as follows:

$$\text{Effect}(C_1) = (1*21) + (-1*37) + \dots + (-1*26) + (1*3) = -86$$

Note that the sign of the effect is not important; only the magnitude is. Thus, C_3 , C_1 and C_2 are the most important controls, in decreasing order, based on the cost of their failure to the enterprise.

PB designs require $2N$ experiments where X is the number of controls under considerations and N is the next multiple of 4 greater than X . This compares favorably to the other approach of full factorial design, which would require 2^X experiments since it requires experiments for all possible input combinations. Not only does a PB design require substantially smaller number of experiments, it has also been verified [60] that it generates results comparable to a full factorial design if the parameters are monotonic with a low degree of interaction, something that holds true in our scenario as explained above. Another alternative design possibility is, one at a time approach (i.e. varying one control while keeping others constant) but it is not practical in our problem statement either because such a design does not measure the impact of interaction among the controls.

3.3 Threat Classification

[11] outlines a general approach to threat assessments. A systematic classification of computer intrusions is provided by [41]. [48] provides a summary of misuse cases. An authoritative taxonomy of threats for web applications is provided by Web Application Security Consortium (WASC) [76]. We have combined these four classifications in the first six rows of Table 3.3.1. Phygo and Furnell [59] have analyzed insider attacks a summary of which is provided in seventh row of Table 3.3.1. In this chapter, these attacks will be used for simulations purposes because our test environment is a web application with the additional concern of insider misuse of information.

Table 3.3.1 Threat Classification used in Experiments Conducted

Class	Attacks Covered
Authentication	Brute Force, Insufficient Authentication, Weak Password Recovery Validation
Authorization	Credential/Session Prediction, Insufficient Authorization, Insufficient Session Expiration, Sesion Fixation
Client-side attacks	Content Spoofing, Cross-site Scripting
Command Execution	Buffer Overflow, Format String Attack, LDAP Injection, OS Commanding, SQL Injection, SSI Injection, XPath Injection
Information Disclosure	Directory Indexing, Information Leakage, Path Traversal, Predictable Resource Location
Logical Attacks	Abuse of Functionality, Denial of Service, Insufficient Anti-automation, Insufficient Process Validation, worms
Insider Attacks	transfer of confidential data, access of prohibited content, access to isolated sub-nets, output redirection

Table 3.4.1. Criteria for Determining Cost of a Security Incident

Cost Type	Impact
Productivity Loss	Productivity cost of impacted employees as well as those involved in handling the incident
Revenue Loss	Direct loss, lost future revenue
Financial Performance	Credit rating, stock price, regulatory fines
Damaged Reputation	customers, suppliers, financial markets, banks, business partners
Other Expenses	software/hardware purchases, travel expenses, contractor costs

3.4 Criteria for Determining Cost of a Security Incident

Cost of security incidents is dependent on a large number of factors such as regulation, industry sector, size of the enterprise etc. Therefore it varies from enterprise to enterprise and criteria around it are generally provided by the corporate risk management department. [1] describes why determining the cost of security incidents is a daunting exercise. In order to generalize the cost of security incidents, a normalized model is provided by Farahmand et al [47] that estimates the cost of an incident on a scale of 1-10 each based on different cost types summarized in Table 3.4.1 This cost model would be used to compute costs for incidents caused by exercising attacks in Table 3.3.1 in our experiments. Kark [38] provides a cost estimation model as well. However, his estimation does not cover all possible costs and offers a range of values and hence, is not suitable for our statistically rigorous approach. Similarly [10, 11] provide models for the narrow domain of publicly announced security breaches only (which is only a very small percentage of all breaches) and so, those are unsuitable for our desire to generate a more comprehensive model.

3.5 Determining PB High and PB Low Values

Center for Internet Security (CIS) [13] is a security industry standards body that provides consensus best practice standards for security configuration. It provides two levels of parameter recommendations for various security products: (a) Level 1: These settings provides the prudent level of minimum due care. System administrators with any level of security knowledge and experience can implement these. (b) Level 2: These settings provide a much higher level of security than level 1 and require expert knowledge to be implemented. Where available, Level 1 settings will be used as the low value and Level 2 settings will be used as the high value for security controls in our PB design matrix.

Many vendors provide pre-built configurations for their products for out of the box (OOB) implementation. For e.g. SUN provides an extensive repository of documentation through its SUN Blueprints library [71]. When corresponding CIS security settings are not available, we have analyzed such repositories to identify the PB low and high values for these security controls.

3.6 Experiment Setup

In the following sections, we will demonstrate through an example how a PB design based method can be used to determine the critical security controls of an enterprise, the configuration of these controls and to conduct the analysis of the impact of changing settings in these controls. The test environment was a simulation of a production web based product data management (PDM) application server. Another reason for choice of this system was the fact that complete risk assessment data was available for it thus enabling the validation of PB design results. The application server resided on a Sun Ultra 6500 server with Solaris 9 as OS, 4 CPU's and 4 GB of RAM. All tests were

Table 3.5.1. PB Values for Experiment Security Controls

Control	P&B Low	P&B High
Router	OOB	CIS Level 2
Firewall	OOB	CIS Level 2
Intrusion Prevention System (IPS)	Vendor Low	Vendor High
Operating System	OOB	CIS Level 2
Host Based Firewall	OOB	Vendor High
Webserver	OOB	CIS Level 2
Servlet Container	OOB Connectors & Apps	Only necessary connectors and apps
Telnet	Telnet	Telnet wrapped in IPSEC
File transfer	FTP	SFTP
SNMP	SNMPv1	SNMPv3
Application Security Settings (J2EE)	Security Manager Disabled	Security Manager Enabled
Data Loss Prevention (DLP)	Vendor Low	Vendor High
Browser	Medium	High
LDAP	OOB	CIS
Log Analyzer	Vendor Low	Vendor High
Database	OOB	CIS Level 2
File Integrity Monitoring (FIM)	OOB	Vendor High

conducted in an isolated network to rule out any impacts to other systems. To make the experiments realistic, a load test imitating the production usage of the system was in progress during the tests.

The set of security controls tested is shown in Table 3.5.1. The PB high and low values of these controls were determined using the resources identified in Section 3.3 (CIS security benchmarks and vendor recommendations). Most commonly, we used out of the box (OOB) security settings as low settings. When determining high values, we used CIS or vendor hardening recommendations. When both, CIS Level 1 and Level 2 were available, we chose Level 2 as the high value.

Security benchmarks simulating the seven classes of attacks identified in Table 3.3.1 were used for attack simulation. A combination of commercial and open source tools was used for this purpose: Foundstone vulnerability scanner, Silk Performer and Nmap. If subjective judgments were required, expert penetration testing was used. If a vulnerability was identified, the exploit was considered complete. Similarly, if an incident was caught by any of the controls and reported within 10 minutes, the effect of that incident was discounted when computing the costs of the attacks successful against that configuration. Because of its deterministic statistical nature, a majority of the design creation and experiment execution was automated.

The cost of successful incidents against a given configuration (a PB design row) was computed using Table 3.4.1 as described in Section 3.3. Each cost type for a given incident was rated on a scale of 1 to 10 (as described in [47]) and all the seven ratings for a given incident were added together to arrive at the total cost of an incident.

3.7 Determining Critical Controls

Improperly chosen security controls for deep dive analysis and testing during risk assessment exercise can create an inaccurate picture of enterprise's risk. This can lead to incorrect decisions on where to direct security investment or inadequate controls thus exposing the enterprise to threats. Thus, using a structured approach to determining critical controls is critical for consistency, good security decision-making and improving the security posture of the enterprise.

In addition to the conduct of risk assessment, there are other reasons why it is important to understand the critical security controls of an enterprise. [3] demonstrates the need to make risk-aware IT investments and identification of critical controls can assist in appropriate direction of security investment (e.g. it might be wise to beef up the investment in critical controls because of their importance to the enterprise's defenses). It

can also help in other cost/benefit analyses such as optimization of resource consumption related to security (discussed in detail in Section 3.2). In case of an impending threat (such as the recent Conflicker virus), there might be a need to quickly turn around risk assessments and understanding which controls are critical significantly speeds up the process. Finally, because of their importance to the enterprise's defenses, special care must be exercised in change and downtime management of these controls.

Below, we present a PB design based approach to identify critical controls. This design empowers the risk assessors by providing a structured approach and guidance on identification of critical controls.

Using the PB low and high values for the configuration of these controls, class of attacks and the enterprise's cost determination criterion for security incidents, a PB design can be constructed. Taking our example configuration from Table 3.5.1 with $X = 17$ controls, our total number of experiments then becomes $N = 40$. Through these 40 experiments, the PB value of each of these 17 controls was computed. Then the controls for each class of attacks were ranked based on the significance of the control (1 = most significant and 17 = least significant). Then, the ranks for each class of attack for each control were summed to determine an overall rank. This overall rank is a reflection of the importance of that control to the enterprise using the cost of breaching that control as criterion.

Results in [66] show the efficacy of our design. Table 3.6.1 shows the overall ranks, sorted in ascending order for the controls based on our PB design. It is clear that only the first five security controls are significant across all security attacks exercised. This conclusion can be drawn because of the large sum of ranks difference between "FIM", the fifth security control and "Application Security Settings", the sixth security control. This indicates that parsimony principle holds true as far as security controls of an enterprise are concerned (i.e. some security controls would be responsible for majority of the protection at the enterprise).

Table 3.6.1. PB Design Results for all Security Controls Sorted by Sum of Ranks

Control	Authentication	Authorization	Client-side attacks	Command Execution	Information Disclosure	Logical Attacks	Insider Attacks	Sum of Ranks
Log Analyzer	1	1	4	4	1	2	2	15
Firewall	2	4	5	2	3	1	4	21
Intrusion Prevention System (IPS)	4	3	2	3	2	3	10	27
Operating System	5	2	6	1	7	4	8	33
FIM	3	6	7	5	5	5	3	34
Application Security Settings	12	13	1	6	4	6	14	56
Data Loss Prevention (DLP)	9	14	11	8	6	8	1	57
Webserver	6	8	8	9	8	7	15	61
LDAP	10	11	12	10	14	9	6	72
Host Based Firewall	13	7	15	11	10	11	5	72
File transfer	17	5	3	14	11	13	9	72
Router	8	15	10	12	9	12	11	77
Database	7	12	14	7	15	17	12	84
Browser	11	16	13	17	16	10	7	90
Telnet	14	9	16	13	12	16	13	93
Servlet Engine	16	10	9	16	13	14	17	95
SNMP	15	17	17	15	17	15	16	112

It is easy to see which control is most important for a given class of attack using cost to the enterprise as the criteria. As an example, firewall seems to be most effective in preventing logical attacks components of which include denial of service (DoS) attacks. This is an expected outcome since stateful firewalls can be configured to prevent DoS. What is surprising though is the importance of “Operating System” security settings in preventing DoS. However, it is not surprising anymore when one analyzes the security settings of Solaris 9 in detail; it does provide TCP and ARP settings that are effective in preventing DoS. The fact that the PB design recognizes the importance of these two in preventing DoS validates our approach to some degree.

“Log Analyzer” appears as the most significant security control. At the outset, this challenges the traditional thinking amongst risk assessors who almost always place a premium on firewalls. However, further analysis of the log analyzer functionality indicates why it might be the most important security control. The log analyzer is the “see-all” security control since most other security controls (including firewalls) feed their logs to it. The value added by the log analyzer is its ability to correlate across all these logs from individual controls and then, smartly identify which events merit further investigation. Thus, because of its see-all status, it is elevated as the most important security control.

As mentioned above, one of the reasons this PDM system was chosen for experimentation was because of the availability of a full risk assessment conducted by experienced risk assessors, which can be used to verify and validate our PB sum of ranks approach. Our analysis of the data from this full risk assessment indicated that our top five controls reconciled with the top five controls identified through this full risk assessment. This validates the PB sum of ranks approach at least to the extent that it can be used to enlighten and focus the activities of the risk assessors upfront thus alleviating the need to conduct a full risk assessment in most circumstances. Also, as compared to tediously manual process of conducting a full risk assessment, the PB design approach

can be substantially automated. This leads to significant savings in time and resources expended. In addition, it improves the security posture of the enterprise in face of impending threats (such as a virus infestation) by enabling risk assessment activities to be carried out expediently.

3.8 Determining Configuration of Controls

The PB design helps in identification of critical controls of an enterprise. This simplifies the process of determining the security configuration considerably since now, only the values of key controls needs to be determined carefully since they are the biggest factors in overall security of the enterprise. This enables sensitivity analyses for each critical security control using ANOVA [65] with gradual change in their values thus leading to the identification of the most effective configuration of those controls. As demonstrated in [65], ANOVA can be effective in determine the configuration of controls when the number of controls under considerations is small (less than 8). While given the importance of the macro security architecture of the enterprise, this recommended configuration still need to be validated by an experienced security architect, his/her task is made considerably easier because now, there is a smaller and prioritized subset of controls with recommended configuration that needs verification.

In a nutshell, following is the approach we recommend fro determining the configuration of the controls:

- (1) Determine the set of security controls that need to be analyzed further.
- (2) Determine PB high and low values for these controls.
- (3) Identify cost computation and threats impacting the enterprise.
- (4) Conduct PB experiments and rank the controls as shown in Section 3.3 to identify the critical controls.

- (5) Iteratively perform sensitivity analysis for each critical control using the ANOVA technique [40].
- (6) Choose the final configuration of the critical controls based on the results of step (5).
- (7) Choose the final configuration of non-critical controls based on vendor recommendations, CIS recommendations, in-house analyses or other appropriate sources.

3.9 Measuring Impact of Security Enhancements

Security executives are keenly desirous of understanding the impact of security enhancements within their enterprise [61]. Some example security enhancements of this kind would be: application of security patches, a review of the application code to determine and fix any security flaws or even, a security awareness training through the enterprise. This understanding extensively drives the decision making with regards to security such as which defenses to solidify or how to allocate security dollars.

However, a survey of literature indicates that existing risk assessment methodologies do not address this issue directly. In the simplest case, it is possible to run threat simulations after the change has been made to determine what its impact was on key security metrics (e.g. virus penetrations, worms, malware etc.). However, identification of all such applicable metrics is a challenge. Also, while these metrics provide some insight into

Table 3.8.1. PB Design Results Post Implementation of Security Code Review Recommendations

Control	Authentication	Authorization	Client-side attacks	Command Execution	Information Disclosure	Logical Attacks	Insider Attacks	Sum of Ranks
Log Analyzer	1	1	3	3	1	2	2	13
Firewall	2	4	4	1	3	1	4	19
Intrusion Prevention System (IPS)	4	3	1	2	2	3	10	25
FIM	3	6	6	4	5	5	3	32
Operating System	5	2	5	5	7	4	8	36
Data Loss Prevention (DLP)	9	14	11	7	6	8	1	56
Webserver	6	8	7	9	8	7	15	60
Application Security Settings	12	13	9	8	4	6	14	66
File transfer	17	5	2	14	11	13	9	71
LDAP	10	11	12	10	14	9	6	72
Host Based Firewall	13	7	15	11	10	11	5	72
Router	8	15	10	12	9	12	11	77
Database	7	12	14	6	15	17	12	83
Browser	11	16	13	17	16	10	7	90
Telnet	14	9	16	13	12	16	13	93
Servlet Engine	16	10	9	16	13	14	17	95
SNMP	15	17	17	15	17	15	16	112

what the effect of the security enhancement was, they do not assist in answering key questions answers to which are required to make smart security decisions. For instance, how did the security enhancement change the security posture of the enterprise? Did some security controls become less important because the security enhancement addressed the problem solved by that security control? Did a security control become more relevant post implementation of the control because of a correlation between the enhancement and that control (e.g. if data from the security enhancement needs to be sent to the security log analyzer, the log analyzer becomes more relevant post enhancement)?

We propose an approach in [66] that can be used to understand the effect of a security enhancement within the enterprise. Our method uses the PB design approach to analyze the effect. Essentially, a comparison of the sum of ranks of the controls in the PB design results before implementing the enhancement and after implementing the enhancement provides a good window into what the impact of the enhancement was.

Our test PDM system underwent an automated code review to determine the security vulnerabilities in the code and these vulnerabilities were fixed. Thus it provided a good means to illustrate this method. Table 3.6.1 represents the PB design values and sum of ranks before the enhancements from the code review were implemented. Table 3.8.1 represents the PB design values and sum of ranks after the vulnerabilities identified through the code review were implemented.

A comparison of the two tables yields to some interesting conclusions. First, that although there has been minor change in the ordering of the critical controls, the set of critical controls (i.e. the five most important controls for the enterprise) remains unchanged. Second, that “Application Security Settings”, and “Operating System Settings” have receded in importance. These results are expected results and in a sense, also validate our PB methodology. This is because the biggest issues that are identified and fixed in web application code reviews tend to be vulnerabilities pertaining to cross site scripting and buffer overflows. Fixing these vulnerabilities will have the most

important impact on “Application Security Settings” and “Operating System Settings” (note that Solaris 9 has security mechanisms in its kernel to prevent buffer overflows). Since these issues have been addressed in code, now they don’t need to be addressed through these two controls.

The power of this approach is evident from the example provided above. It quantifies the effect that a security control enhancement has on the controls. This guides security architects in determining how the security paradigm has changed as a result of a security enhancement thus increasing his/her efficacy. It also assists security executives in understanding how security investments need to be prioritized post implementation of a security enhancement.

3.10 Summary

Risk assessments are critical for most enterprises to manage security risks as well satisfy various compliance requirements. However, currently prevalent risk assessment methodologies, while provide a structured platform, do not provide granular guidance on conduct of some critical steps of the risk assessment exercise. This chapter clarifies the need for methodological improvement in the way risk assessments are currently conducted. Specifically, they do not address the problem of critical control subsetting (i.e. how to identify critical controls), how to determine the configuration of these controls and finally, how to determine the impact of making security enhancements.

This chapter proposes improvement in risk assessment methodology by providing three statistical design of experiments based methods. The first method addresses the problem of critical controls selection through a P&B design based on the threats faced by it and its cost determination criteria. These key controls need to be managed carefully because of their importance to enterprise’s defenses. Once critical controls have been identified, the second method can be used to determine the configuration of these controls. Lastly, this

chapter proposes a novel approach for determining the impact of changing control configuration. This is accomplished through usage of PB design to rank the controls and determine the change in rankings before and after the change has been made.

In conclusion, this chapter introduces statistical rigor to improve quality and efficiency of the risk assessment process. Adopting methods in this chapter can lead to a truer picture of how an enterprise needs to manage its controls so that the risks faced by it are within its risk appetite.

Chapter 4

Dynamic Security Configuration Management

One essential aspect of being able to manage the information security risk to the enterprise is configuring security controls appropriately to ensure that the organization is protected against the threats impacting it. However, despite this critical need, there is a significant opportunity in current approaches that are used for this purpose. They are initially configured during the installation phase and then changed only on an event driven basis. These events could be things like an incident, or observation from logs or recommendations from a risk assessment exercise. There are significant issues with this approach: these changes are ad-hoc and either happen *after the fact* (i.e. the loss to the enterprise has already happened at that point) or are not *dynamic* in nature (it makes sense to manage security configuration as soon as the security controls start sensing that the nature of threats around it has started changing).

In this chapter, we propose a novel statistical design of experiments based security architecture for ongoing security analysis and generating security control configuration change recommendations based on the cost criteria important to the enterprise and the changing nature of threats. For the purposes of this chapter, our reference to controls means “technical controls” only, which are devices, protocols and tools used to protect the enterprise such as firewalls etc.

Rest of this chapter is organized as follows: Section 4.1 describes the STARTS architecture in detail; Section 4.2 describes the experiment setup and results and Section 4.3 summarizes this chapter.

4.1 STARTS Architecture

Following is a description of STARTS (Statistically Rigorous Techniques in Security) architecture and its components. It enables better risk management through ongoing statistical analysis of incoming threats based on costs to the enterprise and generating recommendations on the basis of that.

4.1.1 Critical Security Controls

Given that the total number of security controls at most enterprises can run into double digits, configuring every control in the analysis environment (described in detail in Section 4.1.2) is an expensive proposition. Therefore STARTS needs a critical security control determination model to help narrow down the security controls that need to be configured in the analysis environment thus optimizing on costs.

In addition, this step needs to be executed only on an as needed basis and the changes made in the STARTS analyzer to ensure that it is producing relevant results: in situations such as addition or removal of security controls or a significant change in the security posture of an enterprise. For our case study of a retail environment, this translates into this step being exercised on an average of four times a year. Finally, this exercise can be substantially automated through security benchmarks and load attack simulation software thus reducing the overhead costs.

We will be using the Plackett-Burman designs to determine critical controls (as described in Chapter 3). In summary, a PB design needs to be constructed with all the controls using the class of attacks pertinent to the enterprise (typically captured in the risk assessment metrics of the enterprise). The costs are determined using the enterprise's cost determination criteria for security incidents. Then, the ranks for each class of attack for each control are summed to determine an overall rank. This overall rank is a reflection of

the importance of that control to the enterprise using the cost of breaching that control as criterion.

4.1.2 STARTS Components

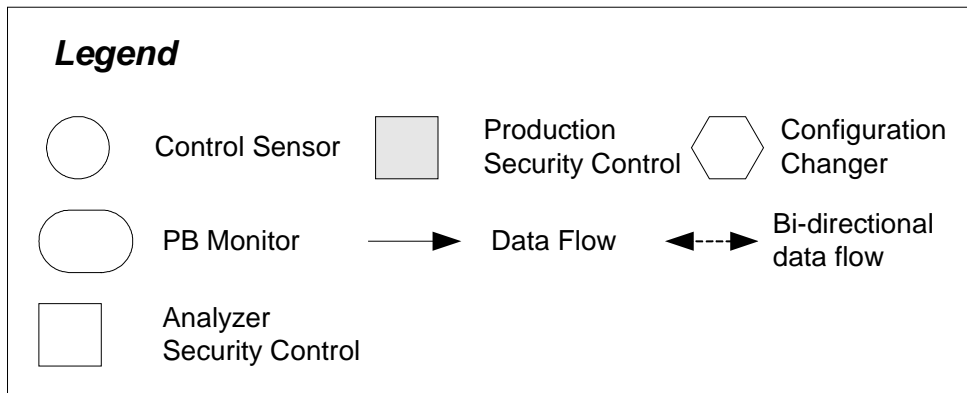
Fig. 4.1.3.1 shows the components of the STARTS architecture. The heart of this architecture is an analyzer that is constructed in its own independent environment distinct from the production system. A regeneration TAP (shown in Fig. 1(a)) that sits after the border router in the DMZ is used to duplicate the network traffic one branch of which is fed into the analyzer and the other branch goes to the production system. This traffic is one-way to protect the enterprise infrastructure from any contamination of the analyzer environment.

Note that many enterprises have honeypot environments already. If so, that same environment can be leveraged as the analyzer environment as well. In fact, STARTS architecture lends an added layer of legitimacy to it (because of the existence of real security controls that attempt to thwart the attacker) thus leading to a higher chance of luring the attacker away from the key infrastructure.


Following is a definition and explanation of the components of the architecture:

Analyzer Security Control

This is the set of critical security controls that was identified through threat analysis and prioritization conducted via PB method as described in Section 3.1 (based on the experiments that we have conducted, the top 1/3rd of the security controls paint a pretty accurate picture of the changes that need to take place to keep the enterprise well protected). As has been shown by Singh & Lilja [65], the parsimony principle applies to the security controls (i.e. a small number of security controls are responsible



(a) A fork in the incoming sends the same network traffic that goes to the production system to the analysis engine as well.



(b) Logical diagram of the components of the STARTS architecture. Notice that the logical organization is independent of the physical location of the control sensors or the controls. Also, components to the left of the thick dotted line are a part of the analyzer.




Fig. 4.1.3.1. The STARTS Architecture

for a majority of the protection offered to the enterprise). Hence, for cost and process optimization reasons, the analyzer functions only with the critical security controls while providing most of the benefit. If cost is not a consideration, the analyzer can be configured with *all* the same security controls as the production system.

Configuration changer

Configuration changer is responsible for changing the security configuration from one PB row to another PB row. The time interval for how often this change happens is determined by the nature of threats to the enterprise and the time it takes for these threats to achieve fruition. As an example, in an eCommerce environment, a denial of service attack is a major concern. A denial of service attack takes a couple of hours to achieve fruition; so a time interval of 45 minutes to 1 hour would suffice.

Note that the PB matrix is significantly smaller now vs. the analysis done in Sections 3.1 and 3.7 because we are configuring the analyzer environment only with critical security controls.

Control Sensors

Control sensors are simple software components that analyze the logs from analyzer security controls and determine the costs. The control sensors for STARTS were written in Perl. Controls sensors are passive components in STARTS i.e. they don't actively solicit or feed data. The cost data generated by control sensors is written to a kernel protected log file that is read by PB monitor.

PB Monitor

PB monitor serves dual roles. It is responsible for gathering cost data from control sensors and generating a new PB matrix every time new cost data becomes available for

an analyzer security control. If PB matrix indicates that a change needs to be made, it generates a recommender alert along with the data for security architect to examine and implement.

4.1.3 Functioning of STARTS

Once the analyzer environment has been created, STARTS uses a dynamic PB matrix to determine if changes need to happen in the production environment. The rows in this PB matrix are analyzer security controls configured as PB low or high. The columns in this PB matrix are the classes of attacks that have been identified as relevant to the enterprise.

At the beginning of an analysis cycle, the configuration changer configures the analyzer security controls according to a given row of the PB matrix. Then the analyzer environment is allowed to function for a predetermined time period based on the volume of the network traffic as well as an analysis of the threats to the enterprise. At the end of this period, three things happen:

- 1) The configuration changer configures the analyzer security controls according to the next row in the PB matrix.
- 2) The control sensors analyze the logs to determine the cost in each security control for the previous time period.
- 3) The PB monitor pulls these costs and ranks the controls using the same approach as described in Sections 3.1 and 3.7.

This cycle is continued on a rotational basis (i.e. when the last row is reached, the process restarts at the first row).

If the threats to the enterprise remained static, the ranking of the security controls in this environment would be the same as what was arrived at when identifying critical controls.

If the rankings change, it is an indication that the nature of incoming threats to the enterprise has changed and that the security controls that achieved a higher degree of criticality need to be configured to a higher level of security. An alert would be generated for the security architect in this scenario who would then use expert judgment to reach a decision on whether to elevate the security configuration as suggested by the PB monitor or not.

In addition, if some controls go down in criticality, that's an indication that threats of that nature have reduced in occurrence indicating that, with expert analysis to back it, there might be an opportunity to lower those defenses thus leading to optimal utilization of system resources dedicated to security. To emphasize though, any lowering of guard because of recommendation from STARTS should happen after analysis and recommendation from multiple security architects to ensure that such a change is appropriate.

4.2 Experiments and Results

This section shows the results from preliminary experiments conducted in a test environment, which was a simulation of a production web based product data management (PDM) application server. The application server resided on a Sun Ultra 6500 server with Solaris 9 as OS, 4 CPU's and 4 GB of RAM. All tests were conducted in an isolated network to rule out any impacts to other system.

We identified a total of $X=17$ controls as shown in Table 3.5.1 [66]. This led to a total of $N=40$ experiments. The PB high and low values of these controls were determined using CIS security benchmarks and vendor recommendations. Most commonly, we used out of the box (OOB) security settings as low settings. When determining high values, we used CIS or vendor hardening recommendations. When both, CIS Level 1 and Level 2 were available, we chose Level 2 as the high value. We used class of attacks described in Table

3.3.1 and cost determination criteria as described in Table 3.4.1 and [47].

Table 3.6.1 shows the overall ranks, sorted in ascending order for the controls based on our PB design. We used top 7 security controls to construct our analyzer environment. This leads to the same PB matrix being used as defined in Table 3.2.1.

Following security controls were selected for the analyzer environment as a result of this exercise:

- Log Analyzer
- Firewall
- Intrusion Prevention System (IPS)
- Operating System
- FIM
- Data Loss Prevention
- Webserver

These controls were configured in the analyzer environment in the same architecture as the production system.

In order to determine the efficacy of the STARTS system, we conducted the following experiments:

- All seven threat benchmarks as identified in Table 3.3.1 were exercised in our isolated test environment using the same configuration of security controls as currently in the production system. Total costs were computed for this setup.
- In the second round of tests, the test environment was configured with the STARTS architecture. Configuration changer rotated the control configuration every 15

minutes to a new PB row. Now the same seven threat benchmarks were exercised here as well. Any recommendations of change from the PB monitor (i.e. observation of change in the rankings of the controls) were always implemented by a security architect. Total costs were computed in this environment as well.

The control sensors were configured to determine costs based on top 20 events in the logs. They were rated on a scale of 1-10. Then log events can be given a cost and added together for a given class of attack to come up with an overall cost; [47] discusses this approach in greater detail. Control sensors were implemented in PERL.

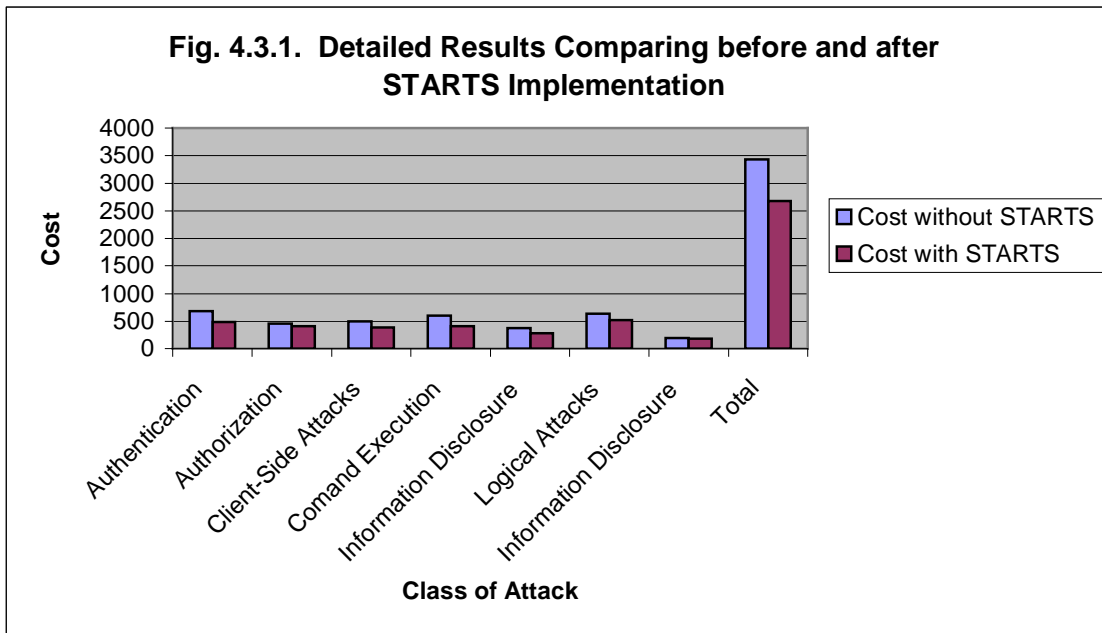
Configuration Changer was setup to change configuration from one PB row to another PB row every 30 minutes. The reason for such a short time interval choice was because we were exercising a sustained attack against the environment over an extended period.

PB monitor was implemented in PERL as well. Given that PB monitor is an active software component, it was run in the privileged mode.

Our experiments showed that when using the STARTS architecture, the cost of the incidents in total was 22% lower than when not using it [67]. More detailed results are shown in the Fig. 4.3.1.

STARTS addresses some key deficiencies in existing security configuration approaches. Specifically, following are some strong points of the STARTS architecture:

- It is dynamic in nature and offer recommendations on configuration changes as soon as it senses any change in the enterprise's threat profile and security posture.



- It can be easily customized to the enterprise because it uses the threats that are relevant to the enterprise as well as the cost determination criteria used by it for risk assessment/management purposes.

STARTS uses the novel approach of multi-level PB design that helps scalability. The first design, which needs to be conducted only on an event driven basis (such as addition of a new security control) is used to identify critical controls which narrows down the analyzer controls thus significantly reducing the enterprise costs to implement additional controls in analyzer environment, computational costs to compute the changes that need to take place as well as the number of experiments that need to be conducted on an ongoing basis in each cycle by the analyzer.

STARTS addresses some key deficiencies in existing security configuration approaches. Specifically, following are some strong points of the STARTS architecture:

- It is dynamic in nature and offer recommendations on configuration changes as soon as it senses any change in the enterprise's threat profile and security posture.

- It can be easily customized to the enterprise because it uses the threats that are relevant to the enterprise as well as the cost determination criteria used by it for risk assessment/management purposes.
- STARTS uses the layered approach of multi-level PB design that helps scalability. The first design, which needs to be conducted only on an event driven basis (such as addition of a new security control) is used to identify critical controls which narrows down the analyzer controls thus significantly reducing the enterprise costs to implement additional controls in analyzer environment, computational costs to compute the changes that need to take place as well as the number of experiments that need to be conducted on an ongoing basis in each cycle by the analyzer.

Since STARTS requires a preliminary analysis to determine critical controls which needs to be repeated on an event driven basis, it does not suit enterprises where changes in enterprise security architecture are happening all the time because than the benefits of STARTS are far outweighed by the costs of conducting the analysis to determine critical security controls on a very frequent basis.

The first implementation of STARTS was done as a prototype using PERL. This initial implementation had the limited objective of validating the architecture as well as generating feedback for the second and more complete prototype. Specifically, as a result of knowledge gleaned from this first prototype as well as research in other security architectures, following enhancements are planned for the second implementation :

- Since new control sensors need to be developed for analyzer environment when new controls are added to it, it is important to have API's for this purpose. STARTS architecture has been enhanced to include these new API's.

- The components of STARTS (specifically the PB monitor) house sensitive information, which can be exploited for malicious purposes. So it is important to ensure that it runs in the privileged mode (such as in the kernel of the OS).
- STARTS must impose minimal overhead on the system it runs to ensure that it is not starving the system of the resources needed to function effectively. Fortunately, STARTS architecture is inherently efficient since it uses a prioritized set of controls only in the analyzer environment rather than the complete set.

4.3 Summary

Current security configuration approaches are either static (i.e. determining the security configuration upfront based on estimates) in nature or are event driven (in response to incidents or risk assessment exercises). Both the cases do not provide desired level of security to the enterprise since risks to their information is very dynamic.

This chapter clarifies the need for methodological improvement in the way enterprises currently configure their security controls. To address this problem, we have proposed STARTS, a statistical design of experiments based architecture. Specifically, a Plackett & Burman (PB) model can be used to determine the critical controls. These critical controls can then be configured in a test bed to which the one-way network traffic can be forked for ongoing analytics through control sensors, which collaborate with each other via a smaller PB matrix housed by PB monitor. The change of analyzer control rankings drives the recommendations provided to the security architecture for production security configuration adjustment.

Our research introduces statistical rigor to dynamically configure the security controls of an enterprise based on how the threats around it are changing. Our experimental results

indicate that our approach can lead to a significant drop in the successful incidents, thus improving the security posture of the enterprise.

Chapter 5

Managing Security Risks Due to Vendors

Most enterprises share data and infrastructure with their vendor service providers. As has happened many times in recent past (an example being the recent exposure of applicant data at CheckFree.com which provides electronic check payment services to several major banks [16]), these shared resources can expose the enterprise to security breaches if the vendor does not have appropriate security controls in place. Consequently, enterprises are realizing that there is a need to extend information security due diligence to vendors as well to ensure the safety of its data and infrastructure. In addition, compliance regimes such as PCI, HIPAA and GLBA are increasingly honing in on this area as well. As a result of this important need, some components of how to exercise this due diligence (such as how to conduct vendor information security maturity assessments) have been defined by industry bodies. However, these components are not in widespread usage yet because of the absence of an end-to-end framework for this purpose as well as a lack of understanding of roles and responsibilities. In this chapter, we present VIAP framework. It provides a structured approach for conducting information security due diligence on vendors. It also establishes roles and responsibilities in an organizational context thus making the task of adaptation simpler.

5.1 Security Risks Introduced by Vendors

The acclaimed management philosopher Peter Drucker stated in 1956 that since the purpose of a business is to create and keep a customer, a business enterprise has only two basic functions: marketing and innovation [15]. Although it seemed farfetched at that time when almost all functions were managed in-house, in today's world, it is

commonplace for most enterprises to focus on their core areas of marketing and innovation while outsourcing other non-core functions to their vendors.

A side-effect of this dependency on vendors for non-core functions however is that it can introduce significant risks to enterprise's information and infrastructure. As an example, a healthcare services provider to the employees of the enterprise will have access to private information such as social security numbers any breach of which can be significantly damaging. Another example is the commonplace situation of offshore vendors having access to corporate infrastructure for project development and support purposes; if not handled appropriately, these accesses can lead to security breaches. Recent security breach of CheckFree.com [16] is a real world example of how corporate data can be put at risk because of sharing of data and infrastructure with the vendors.

As a result, executive management, particularly CSO's, CISO's and Compliance Officers have started focusing on ensuring the integrity of corporate data and systems that vendors have access to. This point is being further driven home by compliance regimes such as GLBA and PCI which have added this dimension of security to their standard audits as well.

Despite this increasingly becoming a critical need, there is a dearth of comprehensive frameworks for managing this risk. Financial Institutions Shared Assessments Program (FISAP) was an initiative started in 2006 to streamline the objective of identifying the risks associated with vendors [24]. FISAP is focused on narrow objective of making the vendor assessment process more efficient; it does not address the larger question of a framework for how to manage risks associated with the vendors. Before FISAP, most vendors had to respond to audits and assessments from multiple financial organizations. FISAP created a single questionnaire that the vendors could complete once and send it to all financial organizations that they serve thus minimizing the overhead associated with audits and assessments. Since 2008, the scope of FISAP has widened to incorporate the needs of non-finance organizations as well and its name has changed to Shared

Assessments to reflect this change. In addition to FISAP, many vendors subject themselves to audits from industry bodies, which they share with their clients to demonstrate their IT Security readiness. Statement on Auditing Standards No. 70 (SAS 70) [56] from American Institute of Certified Public Accountants (AICPA) and ISO 27005 [35] audits are commonly used for this purpose.

While the related work discussed above does not provide a comprehensive framework for managing security risks related to vendors, it does provide several foundational elements that we will use in VIAP. Specifically, we will use ISO 27001 to define a sample vendor information security standard. Our framework is independent of the approach used to conduct security assessment on the vendors; either FISAP or any audit reports acceptable to the enterprise will suffice for the purpose.

In this chapter, we propose Vendor Information Assurance Program (VIAP), an end-to-end framework for this purpose. This framework has been arrived at through interviews with experts in different industry segments, examination of industry best practices, the authors' experience in vendor security due diligence for multiple enterprises and the use of requirements engineering techniques. It incorporates existing components from industry bodies (such as Shared Assessments that are used for vendor information security assessments and ISO 27005). In addition, it defines organizational roles and responsibilities and workflows for how this risk needs to be managed.

Rest of this chapter is organized as follows: Section 5.2 describes current state of the art and how it would be used in our framework; Section 5.3 provides VIAP overview; Sections 5.4, 5.5 and 5.6 describe VIAP in detail; Section 5.7 summarizes this chapter.

5.2 VIAP Framework Overview

This section provides an overview of the VIAP framework. We begin by describing the key stakeholder teams and their role as it is crucial to understanding and implementing the framework. General classification of key stakeholders is provided in [77] which has been adapted for VIAP below:

- **Executive Management:** [26] suggests that executive management of the enterprise ultimately owns the responsibility of maintaining oversight over the risks introduced by the vendors. It exercises this oversight by declaring its intent through security policy statements [57] and ensuring compliance with it through reporting.
- **Vendor Management:** The role of vendor management is to manage request for proposals (RFP's), Vendor management also *manages* information security maturity assessment of the vendors and the remediation of any findings. Vendor management is the most critical stakeholder in VIAP.
- **Information Security:** Information security team is a service provider to vendor management. It *conducts* information security maturity assessment of the vendors identified by the vendor management based on enterprise standards. It also provides consulting services such as how to remediate any findings to the vendor management.
- **Internal Audit:** The job of internal audit is to ensure the sanctity of the VIAP framework implementation. It accomplishes this through random audits of the vendors who have been examined. Internal audit serves as eyes and ears of the executive management.
- **Legal:** Legal department owns the enterprise contracts with the vendors. As such, it consults with vendor management and the information security departments on best

security terms to negotiate into the contracts, the SLA's to establish in case of incidents etc. and any audit requirements.

As is evident from the key stakeholder list above, people impacted by or involved in vendor relationships permeate the entire organization. This makes the task of information security due diligence complex. VIAP introduces simplicity to this issue by partitioning this area into three non-overlapping key components:

- **Vendor Security Governance:** Security awareness is one of the most important and one of the most overlooked [34] means to ensure that the enterprise is better protected. In the context of vendor management, it means that the executive management should have clearly articulated views on vendor security management, Information Security department should have clearly published security standards and all stakeholders should be aware of these and apply them as applicable.
- **Due Diligence on Vendor Contracts:** Enterprises get stymied in their attempts to negotiate changes in vendor information security behavior (such as incident management or audit requirements are met) because these come with extra cost for the vendor and hence, their reluctance to agree to and adhere to any changes. Therefore, information security needs of the enterprise must be negotiated in the contracts.
- **Measuring Compliance with Standards:** This audit is needed to ensure that vendors are complying with the standards that the enterprise has specified for them.

Section 5.3, 5.4 and 5.5 talk about these areas in greater detail.

5.3 Vendor Security Governance

Governance is defined as “the set of responsibilities and practices exercised by the board and executive management with the goal of providing strategic direction, ensuring that the objectives are achieved, ascertaining that the risks are managed appropriately and verifying that the enterprise’s resources are used responsibly” [33]. Policies and measuring compliance with it are the means by which executive management provides governance. *Policies* are high-level statements of management intent, expectations and direction [34].

To ensure governance of vendor security, a vendor security management policy must exist within the enterprise. Figure 5.3.1 shows a sample of such a policy. Significant components of this sample have been adopted from [62] and [73]. This can be used to create a custom policy if one doesn’t exist already.

A policy is typically a high level articulation of management’s intent. As such, it does not provide more granular direction and measurable metrics, which would make the task of adhering to it easier for rest of the enterprise. [25] demonstrates the effective way of writing security policies. *Standards* are used for this purpose. A standard is refinement of the policy to a more granular level and provides the requirements that need to be met for adherence to the policy. Figure 5.3.2 shows a sample vendor information security standard. It is based on the controls and control objectives provided by ISO 27001 [36]. This standard can be used as a starting point if one doesn’t exist already for the enterprise.

Note that just the creation of policy and standard is not going to be sufficient unless it is followed up by extensive propagation through the enterprise. This needs to be accomplished through training. Our recommendation is to make it mandatory for all key stakeholders.

Vendor Security Management Policy

1.0 Purpose

Executive management's intentions for publishing a vendor security management policy are to ensure integrity of the shared data and infrastructure. Executive management is committed to protecting <<Company>>'s customers, employees, partners and <Company> from risks due to vendor relationships.

2.0 Stakeholders

Following is a list of key stakeholders impacted by this policy:

- Executive Management
- Vendor Management Group
- Information Security Group
- Internal Audit and Compliance
- Business Process Owners

4.0 Policy

4.1 Vendors must comply with all applicable <Company> policies, practice standards and agreements, including, but not limited to:

- Privacy Policies
- Security Policies
- Auditing Policies
- Acceptable Use Policies
- Any applicable industry and government regulation

4.2 Vendors management must document:

- The <Company> data and its classification that the vendors have access to.
- The <Company> infrastructure that the vendors have access to.
- <Company> data and infrastructure access can be used to support the business needs of <Company> only and for no other purpose.

4.3 In case of termination of a relationship with a vendor, following must be followed:

- The vendor must surrender all corporate data to <Company Name>
- All vendor access to the enterprise infrastructure must be terminated with immediate effect

.

5.0 Policy Exceptions

- This policy does not apply to the vendors who have no access to <Company> data or infrastructure.

.

6.0 Enforcement

Any vendor found to have violated this policy may be subject to termination.

7.0 Revision History

Fig. 5.3.1. Sample Vendor Information Security Policy

Information Security Policy		
a	Information security policy and its maintenance	The vendor has a security policy that has been published to all employees; it should be reviewed and approved by management on a regular basis to ensure currency.
Organization of Information Security		
a	Security requirements	Security requirements will be identified for the vendor and implemented before any grant of access to the vendor to enterprise data or infrastructure.
b	Contracts	The service contracts with the vendor should explicitly document security requirements, SLA's in case of security breach and audit requirements.
Asset Management		
a	Inventory	All assets shared with the vendors shall be clearly identified and inventoried.
b	Ownership	All shared assets shall have clearly designated owners.
c	Acceptable use	Acceptable use of information and assets shared with the vendor should be documented through an acceptable use policy and implemented.
d	Information Classification	All <company> data that the vendor has access to should be classified according to <company> data classification scheme.
Human Resources Security		
a	Screening	Background checks should be conducted on any vendor employees or contractors who have access to <Company> data or infrastructure.
b	Terms and conditions of employment	As a part of their contractual obligations, vendor employees and contractors supporting <Company> shall agree to abide by the vendor's information security obligation to <Company>.
c	Termination	Upon termination of any vendor employee or contractor working on <Company> projects, their logical and physical access to all <Company> systems should be revoked immediately and all <Company> assets should be returned.
Physical and Environmental Security		
a	Physical security controls	The vendor should implement physical security controls on <Company> information and infrastructure commensurate to the risk.
b	External and environmental threats	The vendor shall design and apply protection against damage from fire, flood, earthquake, explosion, civil unrest, and other forms of natural or man-made disasters.
c	Utility infrastructure	<Company> infrastructure and data shall be protected from utility disruptions such as internet loss, power outage, water supply loss etc.
d	Secure disposal	Prior to any equipment retirement, all <Company> data must be scrubbed from it by the vendor.
Communications and Operations Management		
a	Change management	The vendor shall control any changes to its information processing facilities and system that impact <company>
b	Segregation of duties	The vendor shall segregate key responsibilities to ensure protection against unauthorized or unintentional modification or misuse of <company>'s assets.
c	Segregation of platforms	The vendor shall segregate test, stage and production systems to prevent unauthorized or unintentional disclosure.
d	Protection against malicious code	To protect the integrity of <company> software and information from malicious code, the vendor shall implement detection, prevention and recovery controls.
e	Network security	The vendor shall adequately manage and control its network connectivity to <company> from threats including information in transit.
f	Removable media	There should be procedures in place for management of removable media, including data loss prevention technology as it pertains to <company> data.
g	Audit logging	The vendor shall maintain audit logs of <company> infrastructure and data usage recording user activities, exceptions, and information security events for an agreed period to assist in future investigations and access control monitoring.
h	Monitoring system use	The vendor shall monitor use of <company> infrastructure and data.

Fig. 5.3.2. Sample Vendor Information Security Standard

Access Controls		
a	Access control policy	The vendor shall adhere to <company> access control policy for <company> infrastructure and data.
b	User registration	The vendor shall have a formal registration and de-registration procedure for granting and revoking access to all <company> information systems and services.
c	Privilege management	The vendor and <company> information security team shall restrict and control allocation of privileges on <company> assets.
d	Password management	The vendor and <company> information security team shall allocate passwords for <company> assets through a formal management process and require a quarterly change of passwords.
e	Review of access rights	Vendor shall review access rights to <company> assets twice a year using a formal process.
f	Unattended equipment	Any unattended <company> equipment shall get screen locked after 15 minutes.
g	Clean desk and screen	A clear desk policy for papers and removable media and a clear screen policy for <company> information processing facilities shall be adopted.
Information Systems Acquisition, Development and Maintenance		
a	Input/Output validation	Any software developed by the vendor for <company> or used by it to access <company> data shall have controls to ensure that input data is correct and appropriate.
b	Source code scanning	Any software developed by the vendor for <company> or used by it to access <company> data should be scanned for vulnerabilities and any applicable vulnerabilities should be fixed prior to production usage.
Information Security Incident Management		
a	Reporting	Any information security events shall be reported by the vendor within 24 hours.
b	Roles and responsibilities	The contract between vendor and <company> shall clearly delineate roles and responsibilities in case of an incident.
c	Collection of evidence	When follow-up to incident involves legal action against offending parties, the vendor shall provide all applicable evidence to enable successful prosecution.
d	Remediation	The vendor shall remediate the cause of an incident as soon as possible and share the results with <company>
Business Continuity Management		
a	Business impact analysis	<Company> vendor management, in partnership with vendor relationship manager, shall ensure that a business impact analysis is in place for all key company vendors.
b	Business continuity plans	<company> vendor management, in partnership with vendor relationship manager and vendor shall ensure that business continuity plans are in place for all key vendors and test these plans on a regular basis proportionate to the risk introduced.
Compliance		
a	Legislative compliance	<company> internal audit team shall identify all applicable compliance regimes that the vendor has to adhere to for <company> to maintain regulatory compliance; the vendor shall adhere to the requirements imposed by these regimes on <company> infrastrucutre
b	<Company> security policy and standards compliance	The vendor shall adhere to <company> security policy and standards when accessing <company> infrastrucutre and data.
c	Intellectual property rights	<company> legal department should ensure that the vendor is contractually obligated to protect <company> intellectual property.
d	Compliance audits	<company> legal department shall ensure that the contracts provide <company> a right to audit the vendor for regulatory and <company> policy and standard's compliance.

Fig. 5.3.2. (Contd.). Sample Vendor Information Security Standard

5.4 Due Diligence on Vendor Contracts

Any after the fact negotiations on security with vendors oftentimes don't meet with much success because changes of such are nature are time consuming and expensive to implement. As a result, it is ideal to remove any such ambiguities by articulating terms clearly in request for proposals (RFP's) for vendor selection as well as contracts with the vendors after they have been selected.

In most enterprise, the legal department is the owner of contracts with the vendors. It should consult with the information security teams to ensure that all information security areas of concern have been implemented in the contract.

Following are the information security areas to consider for RFP's and contracts:

- Incident management: Terms regarding roles, responsibilities, ownership and SLA's should be clearly documented. It is important to include notification time frame in case an incident is detected as well as expectations on how soon the situation would be remediated. Also liability from the incident such as loss to the enterprise and any liability to the enterprise customers/users should be clearly negotiated as well.
- Audit requirements: Many enterprises are required to audit their vendors by compliance regimes to ensure that their data and infrastructure is safe. For example, GLBA specifically requires management of outsourcing risk. Similarly PCI requires vendor testing for security. Even in absence of specific compliance regimes, it is important for the enterprise to do at least some audit of its vendors for this purpose. Many vendors are very careful about the amount of access they will provide to outside parties for audit purposes. Therefore it is critical to negotiate specific compliance regimes and their applicable requirements upfront into the contracts. In addition, the right to audit should be negotiated as well if applicable (in some cases,

such as vendors not handling any confidential data, it might not be appropriate to make this a must have point during discussions).

5.5 Measuring Vendor Compliance

5.5.1 Vendor Risk Tiering

The parsimony principle states that a small number of variables are responsible for majority of the effect on the result [72]. This principle is also known as ABC analysis in Operations Management [75]. Applying that principle to our paradigm, it makes sense then to identify the most critical vendors to the enterprise who are introducing the most risk and hence, a majority of the effort of audit and remediation can be directed their way. Hence, vendors providing services to the enterprise must be analyzed for how critical the vendor is to the enterprise from the perspective of information security and assigned a risk tier.

1	Is the vendor onsite or offsite?
2	Location country of the vendor
3	Type of services provided
4	Vendor network connectivity to the enterprise
5	Classification of enterprise data accessible to the vendor
6	Data transfer technology (I.e. FTP, batch, online etc.) as well as frequency
7	Duration and financial terms of enterprise contract with the vendor
8	Business Impact Analysis (BIA)

The number and type of risk tiers can be customized by the enterprise. We recommend the following three risk tiers, which will meet the needs of the most enterprises:

- **Critical:** These vendors are strategically important to the company because of the services they provide and/or the data and enterprise infrastructure they have access to. Any security breach or downtime of these vendors introduces a serious risk to the day-to-day operations of the enterprise. An example would be a bill-pay service provider to a banking enterprise.
- **Moderate:** Vendors in the moderate tier are important but don't introduce as much risk to the enterprise. In case of security breach or downtime, they cause significant inconvenience but can be recovered from and/or replaced. An example would be vendors providing contracting services to the company.
- **Low:** These are the vendors who don't have access to critical enterprise data or infrastructure. Any security breach and/or downtime does not impact the enterprise. These are the vendors that can be most easily replaced. Table 5.5.1.1 provides some criteria that can be used to evaluate the vendor to determine which tier they belong to. These criteria can also be used with a scoring model to make the task of classification simpler.

5.5.2 Vendor Compliance Audit

Vendors need to be audited by the enterprise to measure compliance with the applicable enterprise security standards, security terms laid out in their contracts as well as for external compliance satisfaction purposes (such as PCI, GLBA or HIPAA). Vendor risk tiering helps target the vendors that need to be audited. Most enterprises should audit at least a sampling of "critical" vendors. The decision to audit "moderate" vendors is

dependent on internal resource availability. No purpose is served in auditing “low” vendors.

VIAP defines the steps involved in vendor compliance audit. The flowchart in Figure 5.5.2.1 shows these steps. Further explanation of key steps is provided below:

1. Vendor Management Sends Questionnaire/Request for Information (RFI): Given that vendors are likely to be geographically very distributed, it is ideal to conduct as much analysis as possible within the enterprise to minimize costs. Sending a questionnaire or RFI to accumulate information serves this purpose. The request could be for any audit documents (such as SAS 70 or ISO 27001) that the vendor has; it could be for vendors to complete a questionnaire that is custom to the enterprise; or it could be for industry standard questionnaire such as the one provided by Shared Assessments [24].
2. Information Security Team Analyzes the Response: Once the vendor has responded to the questionnaire, vendor management passes it on to the information security (IS) team for analysis. IS compares vendors the response against expected responses that it creates on the basis of enterprise’s security needs (such as enterprise security standards, security needs of the enterprise or security language included in the contract with the vendor). Based on this analysis, IS might recommend continuation of that analysis through an onsite visit (which is coordinated by vendor management); or it might identify remediation items or it might close the vendor audit as satisfactory. Note that the truthfulness of the response is ensured by the contextual knowledge of the Information Security team as well as the mechanism of “Onsite Visit” as explained in (3).
3. Onsite Visit: Based on preliminary analysis of vendor response, the information security team might recommend that an onsite is warranted. This could be because some items can be analyzed only onsite (such as adherence to clean desk requirement or physical security) or because for confidentiality reasons, the vendor might agree to share further details only at its own site instead of handing it over.

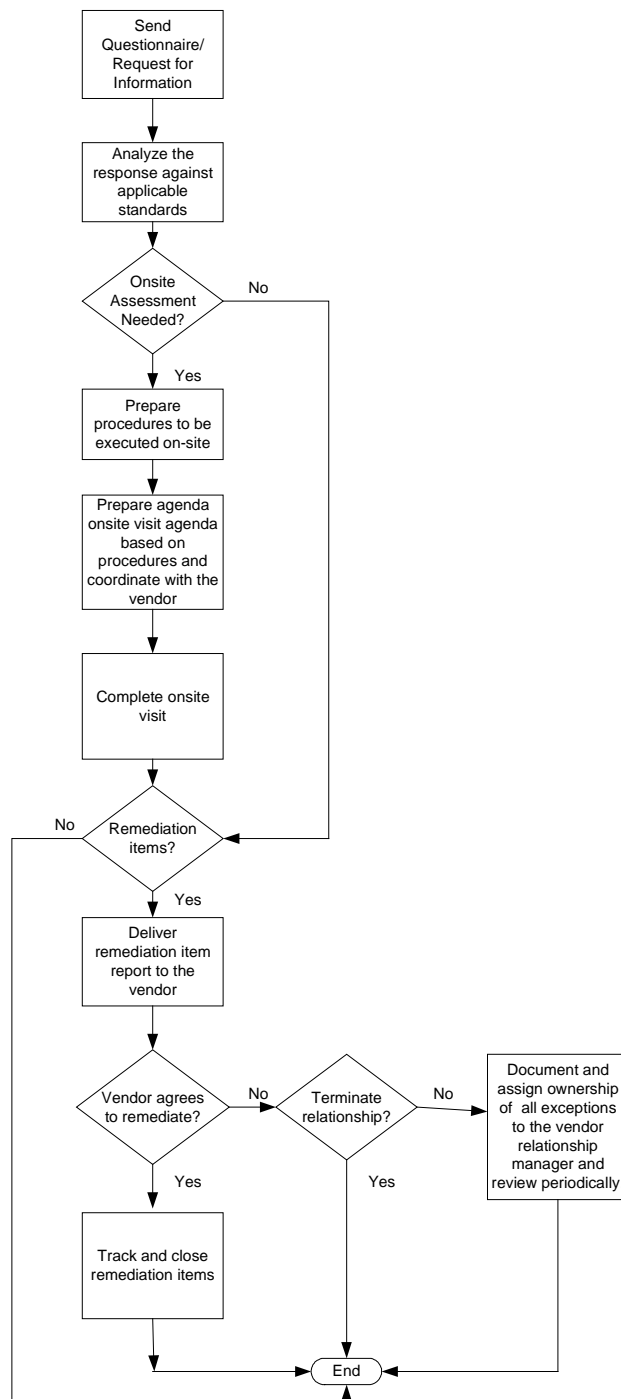


Fig. 5.5.2.1. Flowchart depicting steps in vendor compliance audit

The onsite visit is coordinated by vendor management with the vendor; the information security due diligence during the visit is done by information security team. The outcomes of the onsite visit could be identification or remediation items or complete satisfaction leading to closure of the audit.

5.5.3 Remediation Management

The last step in vendor compliance audit is remediation management of the items identified by IS team. This step is primarily owned by vendor management with security consulting provided by the IS team. Following are the key steps involved in this process:

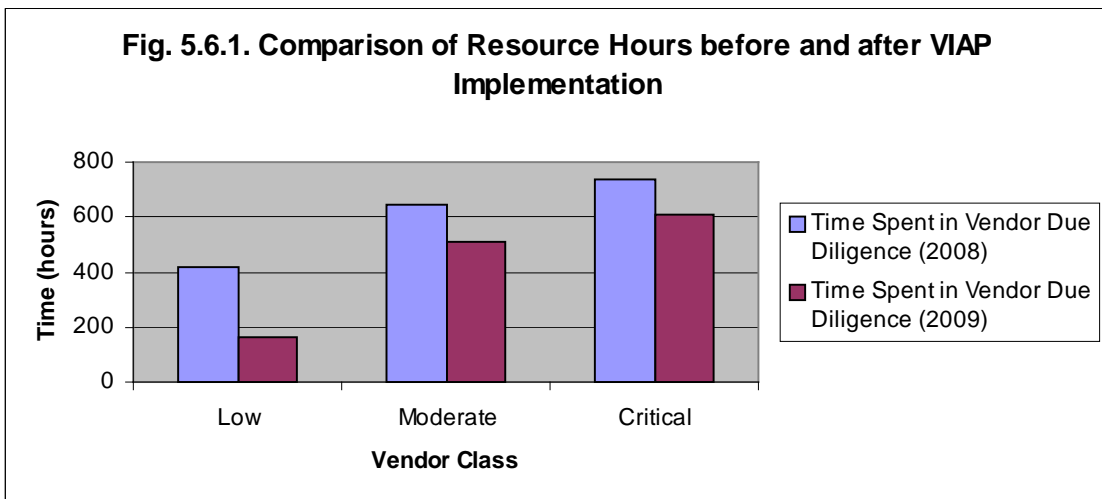
1. **Present Remediation Items to the Vendor:** Once remediation items have been identified, they should be presented to the vendor and then, provide an opportunity for the vendor to respond. IS team acts as a key partner throughout. Vendor might come back with additional detail that might resolve the finding; or it might agree to implement steps to resolve the finding; it might also refuse to do anything about it.
2. **Handling Vendor Refusal to Address the Finding:** If the vendor refuses to resolve the finding, it should be routed through a standards exception process. That entails analysis of the finding to quantitatively or qualitatively evaluate the risk associated with it. This analysis is then presented to the vendor relationship owner (VRO) within the enterprise. The VRO might reach a conclusion that the risk is great enough that the relationship with the vendor needs to be terminated; or in the alternative he/she might accept the risk and continue the relationship. If the VRO accepts the risk, he/she should be documented as the owner of the risk. Based on the magnitude of the risk, VM might want to revisit this risk on a periodic basis.
3. **Managing Remediation:** If the vendor agrees to remediate, VM should work with the vendor on timeframes and verification of remediation post confirmation from the

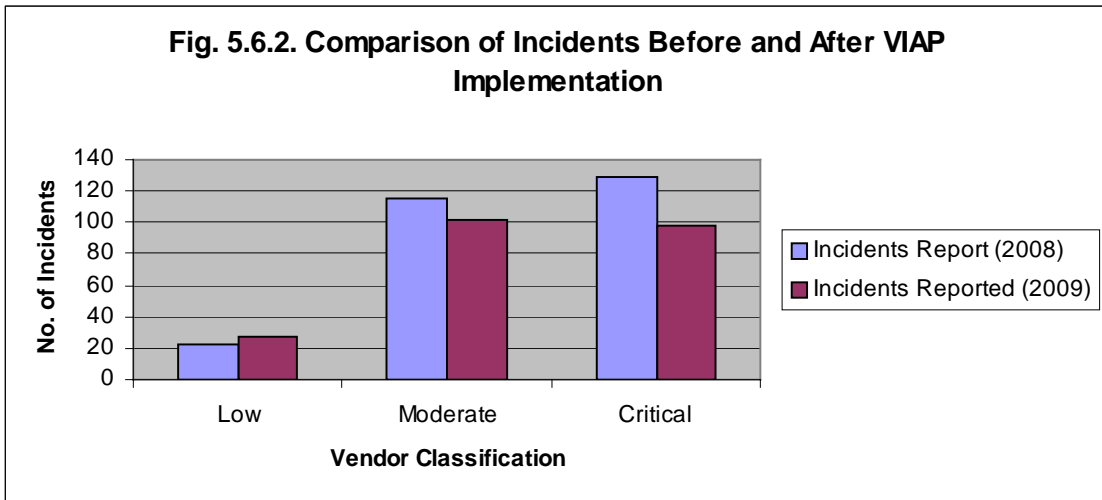
vendor that it has been accomplished. Verification might involve participation of IS teams as well. Once verification has been confirmed, the finding against the vendor can be closed.

5.6 Implementation Results

We have implemented VIAP framework (along with SIG questionnaire to analyze the information security maturity of the vendor) at a large enterprise. Currently, we are six months into the implementation of this framework. [68] documents the results of these experiments in detail. Following charts show those results.

Note that the big difference in resource hours spent for low vendors comes from the fact that after the vendors have been prioritized, the enterprise can choose to spend its time on the more important classes of vendors (moderate and critical) rather than the low vendors.





Following are the observations and results from our experiments:

- 1) Vendor risk tiering was enormously valuable to the enterprise. It helped them identify the key vendors in terms of risk introduced to the enterprise. Given the limitation of resources, this helped senior executives in making smarter decisions on where to spend most time in remediation. As a result, for the first year, the focus was on tier 1 and tier 2 vendors only.
- 2) VIAP clearly outlines roles and responsibilities. This led to better responsiveness as well as improvement in quality of due diligence. In addition, it also led to elimination of areas of overlap. Discounting the one time startup costs, this led to significant optimizations (a total of 17 % resource hours savings were realized).
- 3) VIAP works best for enterprises which have a significant risk exposure from their vendors and which have established vendor management and information security departments. It is an ideal replacement for any ad-hoc solutions for this purpose. For other enterprises (such as small companies or those companies where risk from vendors is low), VIAP might not be the most cost effective solution given its need for

resources.

- 4) Interviews conducted after introduction of Vendor Security Management Policy and Standards indicated that the compliance to it was low. [32] explains that effective training is an excellent means to further compliance to security policies. That led to a mandatory training with certification for all employees involved in the vendor management program. The compliance to the policy increased significantly for the employees who had been through the training.
- 5) The number of remediation items increased significantly as the program got underway. Every remediation item had a workflow associated with it as well. As a result, the remediation component of the framework became very difficult to manage. That led to a realization that implementing a governance, risk management and compliance (GRC) solution with pre-encoded controls mapped to SIG can lead to better efficiencies. The enterprise is in the midst of implementing such a solution.

5.7 Summary

In today's business environment, enterprises are increasingly reliant on vendors to perform critical functions thus leading to sharing of enterprise data and infrastructure with them. As a result, it is imperative that appropriate due diligence be conducted to protect the enterprise against the resulting risk exposure. Currently, this area is either overlooked or there are methodological deficiencies in the way this activity is conducted. This chapter clarifies the need for methodological improvement in the way enterprises conduct this due diligence.

To address this critical need, we have proposed VIAP, a framework for managing the risks associated with vendor relationships. VIAP provides a comprehensive model incorporating policies, standards, contracts due diligence and vendor security assessment.

In addition, it identifies key stakeholders and their roles and responsibilities in managing vendor security risks. VIAP can help meet the vendor security needs identified in key compliance regiments such as GLBA, PCI and HIPAA.

Our future work is focused on creating a simpler and less intense version of VIAP for those enterprises where vendor risk exposure is not significant. A simpler, quick risk assessment framework might lead to a better risk/reward equation in those cases.

Chapter 6

Governance, Risk Management and Compliance

The importance of information to modern enterprises cannot be overstated. As a consequence, governance, risk management and compliance (GRC) issues around information have become central to organizational strategies. Investment in these areas has been increasing steadily topping \$32B in 2008, a growth of 7.4 % over 2007 [30].

There is another driver for having appropriate GRC solutions within the enterprise. The enterprise risk management approaches discussed in Section 2.1, our solution to the gaps identified in Chapter 3 and 4, vendor risk management approach discussed in Section 2.2 and the corresponding framework described in Chapter 5, while improving risk management problem create another: how to manage the enormous amount of data that gets generated as a result of conducting the activities described in these approaches. [64] provides examples of the types of data as well as analysis that needs to be conducted on those:

- Issues identified at vendor sites that need to be tracked, reported and addressed.
- Mapping findings from conducting risk assessments to compliance regimes.
- Risk log
- Prioritization of risk items
- Security control configuration
- Measuring effectiveness of risk management program

[22] and [23] also describe a management perspective of how information security risk to an enterprise should be managed which entails comprehensive reporting and oversight.

So it is evident that managing such voluminous, interlinked data is a complex undertaking.

GRC solutions is a breed of solutions that provide a single, federated framework that integrates organizational processes and tools supporting those processes for the purpose of defining, maintaining and monitoring governance, risk and compliance. An appropriately chosen GRC platform can lead to reduced complexities and increased efficiencies. Selecting a GRC platform is a complex endeavor, though, and requires extensive collaboration between business, IT, compliance and audit. A GRC solution can either be developed in-house to ensure that the custom needs of the enterprise are met or a vendor solution can be selected for this purpose. If a vendor solution is chosen, it is a complex task because of the fact that this space is populated with a large number of competing products; AXENTIS, MetricStream, OpenPages, Paisley, Modulo and Archer being but just a few examples. Thus it becomes imperative that the platform selection should be done intelligently to ensure positive return on investment (ROI) [27, 29].

Selecting an appropriate GRC solution needs a substantial investment of time and effort though in addition to the capital investment required for purchasing and maintaining the platform. This area remains un-addressed in current research. To address this problem, we propose a selection criteria and methodology in this chapter.

The following sections provide a comprehensive set of criteria that can be used to evaluate and select a GRC platform for an enterprise. The criteria presented below have been determined through interviews with experts in different industry segments, examination of industry best practices, the authors' experience in evaluating GRC platforms for multiple enterprises and the use of requirements engineering techniques. These criteria can be used as building blocks to which the unique requirements of the organization can be added to arrive at a complete set of requirements that need to be considered. While they are being defined here for GRC platforms in totality, these can

easily be adapted for toolsets addressing individual areas of governance, risk or compliance. The criteria are structured in three major sections: general considerations, functional requirements and non-functional requirements. We will further cement how to use to use these criteria to arrive at a decision through a scoring model and a case study.

6.1 General Considerations

These criteria are general in nature and applicable to all enterprises irrespective of regulations applicable to them, their size or the business sector in which they operate. These are must haves and hence, are generally used for *exclusionary* purposes i.e. to narrow the field of proposals that would be considered. Figure 6.1.1 summarizes the parameters and artifacts that can be used to evaluate vendors against these criteria.

- **Cost:** GRC solutions can vary significantly in costs. While considering costs, it is important to consider the total cost of ownership (TCO). Some important TCO components are hardware, implementation and consulting fees, training, customization, maintenance, security and operational costs. Also, this is a useful metric to have for ROI calculations.

Figure 6.1.1 - Evaluating general considerations	
General Considerations	Example evaluation parameters and artifacts
Cost	software, hardware, licensing, training, customization, consulting, maintenance, security, operations
Vendor Reputation	references, installed base, financial viability (market capitalization, financial results, annual reports
Product Strategy and Vision	product roadmap, R&D headcount, R&D budget

- **Vendor Reputation:** With the growing popularity and demand of GRC platforms, a significant number of vendors have jumped into this space. In addition to there being a surfeit of genuine vendors, the picture is further clouded by some vendors who market GRC solutions which are nothing but thinly disguised versions of their

existing product suite targeting a different space. As competition heats up and market forces weed out weaker players, only the stronger players would survive. Hence, it is important not to get stuck with a solution that might become unsupported in the future either because the vendor has ceased to exist or because it has exited this space. This can be accomplished through a thorough appraisal of vendor's installed base, references and financial viability.

- **Product Scope, Strategy and Vision:** Threats and vulnerabilities are ever changing. The recent financial meltdown is leading to a change in regulatory landscape as well. All of this is a stark reminder that GRC is an ongoing process that might require an expansion of scope. Another driver for this is the fact that many countries are still working towards maturing their regulations and compliance regimes, J-SOX [74] being one such example. Finally, as organizations enter new market segments, they have to adapt to GRC requirements in that space. All of these factors mean that it is important to examine the product scope, strategy and vision to make sure that the vendor has a long-term view of its product offering and has mechanisms to adapt and expand as the landscape changes. Product road map and R&D strength (measured in terms of R&D head count and investment) are some ways to further this examination.

6.2 Functional Requirements

Functional requirements are used to define the behavior of the target software including features and capabilities that determine what a system is supposed to accomplish. In the sections below, we define high-level requirements each of the three principal components, Governance, Risk Management and Compliance as well as for other general functionality.

- **Governance:** ITGI defines governance as “the set of responsibilities and practices exercised by the board and executive management with the goal of providing strategic direction, ensuring that the objectives are achieved, ascertaining that the risks are managed appropriately and verifying that the enterprise’s resources are being used responsibly” [33]. In light of this definition, it is clear that the governance component of the GRC platform must be evaluated for the requirements presented in Figure 6.2.1.
- **Risk Management:** Risk management is the activity directed towards assessing, mitigating (to an acceptable level) and monitoring of risk. The principle goal of an organization’s risk management process should be to protect the organization and its ability to perform their mission, not just its IT assets [70]. Figure 6.2.2 presents the high-level requirements for risk management.
- **Compliance:** Compliance is an increasingly complex task given global footprints of organizations, increase in regulatory environment (which is likely to become even more stringent given the opportunities exposed by the current economic crises) and local regulations. Figure 6.2.3 presents the requirements to ensure that these needs are supported by the GRC platform.
- **Vendor Oversight:** Regulators are increasingly focused on the personally identifiable data (PII) and how organizations manage it with their vendors who have access to this data. For example, healthcare providers to most organizations will have access to PII. They are requiring organizational due diligence to ensure that their vendors have mature information security practices to protect their data. GRC platforms should facilitate this effort. Support for Shared Assessments, the industry standard for determining the maturity of information security practices at a vendor, is one way a GRC platform could demonstrate its strength.

Fig. 6.2.1 - Governance Requirements	
Requirement	Explanation
Business Alignment	Facilitate alignment of governance with organization's business objectives.
Policy, Standard and Procedure Management	Policies are the medium through which management communicates its direction and intent. Standards and procedures are the vehicles used to implement policies across the organization. Therefore, the GRC platform must support the development, maintenance and communication of these.
Oversight	Enable executive management oversight through appropriate reporting mechanisms such as a security and/or compliance dashboard.
Decision Support	Provide cost/benefit and other data to the executive management for decision-making purposes (e.g. risk data can be used to determine the economics as well as justification of security investment).

Figure 6.2.2 - Risk Management Requirements	
Requirement	Explanation
Risk Baseline	It should facilitate development of the risk baseline based on an organization's risk appetite.
End-to-End Risk Management	Risk management is a <i>continual</i> process. It should begin at the conception stage, should be considered throughout the software development lifecycle (SDLC) and end only when the system is retired. The GRC platform must support this ongoing management of risk.
Adaptability	It must be <i>adaptive</i> . Since an organization's risk profile, threats and vulnerabilities change frequently, it is important that risk management should be adaptive to these changes.
Consistency	It must provide <i>consistency</i> i.e. different areas of the same organization should manage their risks in a consistent fashion. This makes the task of risk consolidation simpler and more manageable.
Metrics	It must facilitate collection of metrics about incidents, vulnerabilities and threats. This data in turn can be used for monitoring losses and assignment of cost-effective controls to remediate or mitigate future losses.

Fig. 6.2.3 - Compliance Requirements	
Requirement	Explanation
Regulatory Intelligence	Reporting on global regulatory changes and connectivity with legal/regulatory databases such as WestLaw and LexisNexis.
Requirements and Controls Library	Authoritative libraries of all applicable compliance driven requirements and associated controls
Correlation	Ability to correlate similar requirements across different compliance regulations for efficiency purposes.
Remediation Management	Ability to track identified remediation measures and their progress
Reporting	Ability to generate reports including ad-hoc reports needed for audits.

- **Workflow:** Given the large number of areas and users involved in the GRC platform, the need to manage and distribute work and the need to monitor its progress through all of those steps, a good workflow engine is essential to the success of a GRC platform.
- **Document Management:** GRC platforms are used for organization and management of an extensive body of documentation. In addition to policies, standards and procedures, they are also used for housing organizational controls, test conducted to verify the robustness of these controls and custom attributes. Therefore strong document management features are essential for it to be successful.

6.3 Non-Functional Requirements

Non-functional requirements are used to define the operation of the system or the environment in which the software should run. Since the spectrum of non-functional requirements is very large, we have narrowed the field down to the requirements that are most applicable to the selection of a GRC platform.

- **Security:** GRC platforms house critical information about the security posture of the enterprise including information about vulnerabilities, risks and data as well as its classification. The consequences of a security breach are great: exploitation of vulnerabilities, damage to credibility, financial loss and legal liability. As such, strong security measures should be provided in the platform to enforce not only protection from external breaches (e.g. through encryption) but also from insider misuse of information by allowing enforcement of the two fundamental principles of security: least privilege (i.e. an individual should have just enough permissions and rights to fulfill their role) and need to know (i.e. an individual should have access to specific information only if it is essential for them to carry out their role).

- **Scalability:** Both the amount and the complexity of information resources within organizations are increasing at an exponential rate. In addition, it might be necessary for organizations to scale their GRC platform for new risks and compliance regimens. This could be necessitated by their foray into new market segments, expansion in their global footprints thus making them subject to local regulations or new regulations coming into existence. Determining scalability requirements appropriately upfront provides flexibility for future growth. Because scalability is based on future needs, it requires a certain amount of prediction and estimation to plan for it. An examination of the strategic business plan of the organization for the next few years might provide this insight.
- **Interface:** In order to achieve maximum efficiency from the GRC platform, it is important that it provide interfaces for integration with enterprise applications used to drive business processes (e.g. integration with identity management system or configuration management database (CMDB)). This will help automate data collection, controls and processes and hence simplify the tasks of analysis, reporting and remediation.
- **Usability:** Usability requirements specify the ease of use of a system. Given that a GRC platform would be used by a broad spectrum of users including business, IT, audit and compliance, it is important that their input is sought in evaluating the usability of any platform under consideration. The five parameters that should be considered for this purpose are: ease of learning (evaluated through training and documentation provided), task efficiency (efficiency of the system for frequent users), ease of remembering, understandability and subjective satisfaction [39].
- **Support:** Supportability deals with the ease of customization to meet the unique needs of the organization, incorporation of new features or enhancements and bug

fixes. A good example of a supportability requirement is that for an organization that has to adhere to PCI, the GRC platform vendor should provide updates when the new versions of PCI get released. Maintenance, updates, consulting services and customization are some areas to consider when evaluating vendors against this dimension.

6.4 Example Selection Process Walk-Through

The criteria presented above can be combined with a weighting mechanism to arrive at a decision on which GRC tool to select. A Case study of how our approach was applied to a large enterprise is presented below [64].

***Organization:** A retail organization is looking to strengthen the governance and risk management of its information. It has been classified as a tier-2 vendor for PCI. In addition, it offers pharmacy services in its stores and hence, has to be compliant with HIPAA as well. It has budgeted TCO of \$750,000 for a five-year period for a GRC solution to manage these efforts. It is not looking to include SOX in the ambit of this GRC tool because it intends to continue leveraging its existing point solution for that. Following is a step-by-step description of how it arrived at a decision using the criteria defined above (Figure 6.4.3 shows the results of these steps).*

1. It created an RFP defining the GRC needs of the organization and invited vendor responses. Based on exclusionary criteria, it narrowed down the vendor choices to A, B and C.
2. It partitioned its stakeholders into *primary* (those who are directly impacted by the platform choice) and *secondary* (those that are intermediaries in the selection process) stakeholders. Its primary stakeholders were: Office of the CISO, IT, Internal Audit and Pharmacy process owners. Its secondary stakeholders were: Vendor Management, Business Continuity Planning (BCP) team and Finance.

3. It identified its other requirements, primarily using the requirements solicitation questionnaire shown in Figure 6.4.1 (shown in Figure 7 under “Other Requirements).
4. It weighted all criteria on a scale of 1 to 5 using Figure 6.4.2 (Note that since this article focused on identifying *essential* requirements in the sections above, most of those would be weighted 3 or more; when unique organizational requirements are added, the spread from 1 to 5 would likely be observed). Figure 6.4.3 reflects the weights along with explanations where the choice of a weight is not obvious.
5. It created a committee drawn from primary and secondary stakeholder teams. For vendors still under consideration, this committee rated them against each requirement on a scale of 0 to 5 using consensus method (some stakeholders chose to recuse themselves on occasions as they were not knowledgeable about the requirement under consideration). Note that a vendor should be disqualified if they have a score of 0 on any criteria rated 3 or above (i.e. any criteria of significant interest to primary stakeholders). They used the following as raw data to arrive at a decision:
 - Vendor Demonstrations
 - White papers, spec sheets and other documentation
 - Data from research organizations like Gartner, Forrester, Burton Group etc.
6. It computed a total weighted score for each vendor. Since the scores of vendor A and vendor B are pretty close to each other, it had those vendors bid against each other to reduce costs and ended up choosing B as a result.

6.5 Summary

Businesses are increasingly relying on GRC platforms to achieve synergies across governance, risk and compliance. In the crowded landscape of GRC platforms, arriving at the right choice for an enterprise is a complex decision. It is imperative that all applicable

criteria should be considered to ensure positive ROI. It is also necessary to make the evaluation process as objective as possible.

The proposed approach in [64] helps facilitate business and IT in understanding the essential criteria to consider when evaluating GRC platforms. In addition, it illustrates how these criteria can be rolled into a scoring model to arrive at an objective decision. This ROI driven approach will improve an organization's ability to select the right GRC platform that fits its need and in turn, will help it manage the complexities associated with governance, risk and compliance efficiently.

Fig. 6.4.1: Requirements Solicitation Questionnaire	
1.	What is your biggest GRC area of concern?
2.	What compliance regulations are applicable to your area?
3.	Have you failed any areas of compliance audits in the past? If so, what were the findings?
4.	What improvements would you like to see in your current mechanism for prioritizing security budget?
5.	How do you rate the effectiveness of your security controls?
6.	What would you like to see in the reports indicating the current status of compliance?
7.	How do you evaluate your risk currently? What are possible areas of improvement?
8.	What are critical threats to your area?
9.	How many times have you experienced these threats in the past 12 months?
10.	What area are you more concerned about: insider abuse or external threat? Please provide specifics.
11.	Have any of your end users expressed dissatisfaction with the extra steps they have to go through because of the security controls?
12.	Do you have a good data classification mechanism?

Fig. 6.4.2 - Criteria Weight Determination	
Stakeholder Interest	Score
1-2 Secondary Stakeholder	1
3 Secondary Stakeholders or more	2
At least one primary stakeholder	3
More than 2 (but not all) primary stakeholders	4
All primary stakeholders	5

Fig. 6.4.3 - Decision Table								
Requirements	Weight (W)	Explanation/Comments	Vendor A		Vendor B		Vendor C	
			Rating R(A)	R(A)*W	Rating R(B)	R(B)*W	Rating R(C)	R(C)*W
Governance								
<i>Business Alignment</i>	5		4.1	20.5	4.7	23.5	2.8	14.0
<i>Policy, Standard & Procedure Management</i>	5		4.7	23.5	3.5	17.5	3.5	17.5
<i>Oversight</i>	4		3.4	13.6	3.7	14.8	4.4	17.6
<i>Decision Support</i>	3	Intention to rely on existing tool set as much as possible	4.1	12.3	3.3	9.9	4.4	13.2
Risk Management								
<i>Acceptable Risk Baseline</i>	4		4.7	18.8	4.8	19.2	3.3	13.2
<i>End-to-End Risk Management</i>	3	Mostly off the shelf software means that managing risk across SDLC is not critical	4.5	13.5	2.1	6.3	4.7	14.1
<i>Adaptability</i>	4		2.1	8.4	3.1	12.4	2.3	9.2
<i>Consistency</i>	5		1.8	9.0	4.3	21.5	2.1	10.5
<i>Metrics</i>	5		4.2	21.0	3.0	15.0	2.9	14.5
Compliance								
<i>Regulatory Intelligence</i>	4		3.3	13.2	4.4	17.6	4.3	17.2
<i>Requirements and Controls Library</i>	5		4.1	20.5	4.0	20.0	3.8	19.0
<i>Correlation</i>	3	Since HIPAA and PCI are mostly non-overlapping, being able to correlate across the two is not critical	3.1	9.3	2.1	8.3	1.9	5.7
<i>Remediation Management</i>	4		4.3	17.2	3.9	15.6	2.8	11.2
<i>Reporting</i>	5		4.3	21.5	4.2	21.0	3.3	16.5
Vendor Oversight	2		2.3	4.6	1.5	3.0	3.5	7.0
Workflow	5		3.9	19.5	5.0	25.0	0.9	4.5
Document Management	5		4.5	22.5	4.1	20.5	4.5	22.5
Security	5		5.0	25.0	5.0	25.0	4.5	22.5
Scalability	2	Does not anticipate a change in its regulatory environment	3.8	7.6	5.0	10.0	4.9	9.8
Interface	4		2.2	8.8	4.1	16.4	3.8	15.2
Usability	5		4.5	22.5	4.3	21.5	4.2	21.0
Support	5		4.1	20.5	1.9	9.5	3.0	15.0
Other Requirements								
<i>Import Existing HIPAA Controls</i>	5		4.0	20.0	2.7	13.5	1.3	6.5
<i>Automatic Evidence Collection</i>	5		4.0	20.0	4.0	20.0	2.9	14.5
<i>Project Management</i>	4		3.9	15.6	4.2	16.8	3.3	13.2
<i>Exceptions Management</i>	4		1.3	5.0	3.7	14.8	3.6	14.4
<i>Fit in existing Infrastructure</i>	3	Hardware is a small part of the overall allocated budget	3.4	10.1	1.9	5.7	1.2	3.6
<i>Support for ISO Guide 73</i>	3	Risk calculation method used in some departments	1.5	4.5	5.0	15.0	4.2	12.6
<i>Background check of vendor consultants</i>	1	Most vendors would comply if selected	4.0	4.0	3.3	3.3	4.8	4.8
<i>Segregation of duties</i>	4	Not a strength of the organization currently	4.1	16.4	3.2	12.8	4.4	17.6
Total				448.85		455.4		398.1
	→	Functional Requirement						
	→	Non-Functional Requirements						
	→	Unique organizational requirements						

Chapter 7

Conclusion

A survey of current literature as well as prevalent risk management practices in enterprise environments indicates that there are some significant limitations in current risk management approaches. Although control selection and management is a crucial part of risk assessment process of these methodologies, no formalized methods exist that help manage these aspects. In addition, the area of managing risks due to vendors of the enterprise as well as a requirements engineering framework for determining an appropriate GRC strategy remain unaddressed as well. Following is a summary of these granular issues in existing risk management approaches:

1. **Desirable Characteristics of a Risk Management Methodology:** A clear articulation of what are the desirable characteristics in a risk management system makes it very difficult to determine whether a particular risk management methodology meets the needs of an enterprise or not.
2. **Determination of Critical/Custom Controls:** Risk assessment is a critical component of risk management exercise. Existing risk assessment approaches are long drawn out exercises that don't suffice for the times when there is a need to conduct risk assessments very rapidly (e.g. in situations where a threat assumes alarming proportions in a short timeframe). Also, they do not provide for mechanisms to customize control sets used for risk assessment depending on the needs and risk appetite of the organization. A security prioritization model that takes into account the criticality of security controls would greatly alleviate these issues.
3. **Determining Configuration of Security Controls:** Current state of the art does not address how security controls of an enterprise should be configured based on the threats faced by the enterprise as well as the cost of fruition of these threats.

4. **Determining the Impact of Security Enhancements:** Security controls are expensive and the funds available to address risks are limited. A means to measure the impact of implementing a security control within the enterprise would greatly further the economics of security decision making.
5. **Dynamic Security Configuration Adjustment:** Configuration of security controls remains fairly static in most enterprises and changes only in response to significant events such as a well publicized threat or a security incident. This ad-hoc approach to managing security configurations is not sufficient to adequately protect the enterprise because threats are ever changing and dynamic in nature.
6. **Managing Risks Due to Vendors:** In today's business environment, enterprises are significantly dependent on vendors for non-core services. However, this also introduces significant security risk to the enterprises because the vendor's have access to its infrastructure and data. None of the existing risk management approaches provide for a framework to address the risks associated with the vendors.
7. **GRC Platform Selection:** In most enterprise, being able to manage the large number of issues identified through the risk management process is a challenging undertaking. This includes defining, maintaining and monitoring of governance, risk and compliance. As a result, most enterprises will either build their own platforms or implement external solutions to manage this complexity. However, there no requirements engineering frameworks exist that will help identify the right criteria and how these criteria need to be balanced with user needs to come up with which solution makes the most sense.

In Chapter 2, we propose a set of criteria outlining desirable characteristics in a risk management methodology. In [64], we address how our governance, risk management and compliance requirements can be used to define the requirements that will map to these desirable characteristics.

In Chapter 3, we propose improvements in risk assessment methodology that address the issues identified in 1, 2 and 3. We proposed a statistical design of experiments based

methodology for this purpose in. This is a novel application of statistical rigor in information security risk management [65, 66]. The first problem is resolved by enabling critical controls selection through a P&B design based on the threats faced by and enterprise and its cost determination criteria. These key controls need to be managed carefully because of their importance to enterprise's defenses. Once critical controls have been identified, ANOVA technique can be used to determine the configuration of these controls. Since these controls are critical controls, their appropriate configuration has the most impact on the risk faced by the enterprise. This addresses the second issue identified above. Finally, we propose that the impact of security enhancements can be determined by determine the change in PB rankings before and after the security enhancements have been implemented. This is a novel approach for determining the impact of changing control configuration. This solves the third problem identified above. Through the use of these statistically rigorous techniques, a truer picture emerges of how an enterprise needs to manage its controls so that the risks faced by it are within its risk appetite.

In chapter 4, we clarify the need for methodological improvement in the way enterprises currently configure their security controls. To address this problem, we have proposed STARTS, a statistical design of experiments based architecture [67]. We use the same model as proposed in Chapter 3 to identify the critical controls. These critical controls can then be configured in a test bed to which the one-way network traffic can be forked for ongoing analytics through control sensors. These controls sensors collaborate with each other via a smaller PB matrix housed by PB monitor. The change in analyzer control rankings drives the recommendations provided to the security architects for production security configuration adjustment. This addresses the fourth issue of dynamic security configuration adjustment.

To address the issue of managing security risks due to vendors, we propose VIAP (Vendor Information Assurance Program) framework in Chapter 5. VIAP is a framework for managing the risks associated with vendor relationships [68]. VIAP provides a comprehensive model incorporating policies, standards, contracts due diligence and vendor security assessment. In addition, it identifies key stakeholders and their roles and

responsibilities in managing vendor security risks. VIAP can help meet the vendor security needs identified in key compliance regiments such as GLBA, PCI and HIPAA.

Finally, to address the issue of GRC platform selection, in Chapter 6 we propose selection criteria and methodology that can be used for this purpose. It is a requirements engineering framework that combines general, functional and non-functional requirements into a scoring model that can be used to arrive at an objective decision [64]. This requirements engineering driven approach improves an organization's ability to select the right GRC platform that fits its need and in turn, helps it manage the complexities associated with governance, risk and compliance efficiently.

Chapter 8

Future Work

This chapter describes some future research areas pertaining to improving information security risk management.

8.1 Dynamic Security Configuration Management

While STARTS design has been proven via experiments in enterprise environment, it can benefit from additional research in the following areas ([67] outlines these additional research opportunities in detail):

- **Changing Security Controls**

STARTS design primarily works in the environments where security controls remain fairly static (because the critical control sensors determined through the first level PB design, which is done only on an as needed basis, drive the analyzer security controls that make configuration change recommendations). The area of how to handle configuration changes in the scenario where security controls are changing frequently remains unaddressed in current research.

- **Control Sensor API's**

Since new control sensors need to be developed for analyzer environment when new controls are added to it, it is important to have API's for this purpose.

- **STARTS GUI**

STARTS user interface needs to be enhanced because despite its strengths, it can be rendered ineffective if its users are not able to easily comprehend and interpret the data produced by it.

- **Deployment and Testing**

The best way to further enhance STARTS architecture is by having people use it. We intend to release our STARTS prototype to selected enterprises first and to the general public later, to allow them to experiment with the architecture, try the system, and provide feedback that allows further improvement in it.

8.2 Vendor Risk Management

In Chapter 5, we have proposed VIAP framework to address information security risk due to vendors. While it would suffice well for enterprises who are exposed to significant risk from their vendors and have established departments in place for vendor management, security and compliance, it may not be the best model for small enterprises where the resources needed by VIAP are not in place [68]. In addition, implementing and rolling out VIAP in such enterprises might be cost prohibitive. For such enterprises, a smaller scale model for vendor information assurance should be developed which addresses most of the risk while managing costs.

8.3 Quantitative Measures for Security Decisions

A survey of information security risk management practices in organizations [61] indicates that there is a shift in perspective for the information risk program from that of a

technical specialty to a business advisory and consultancy. The goal then becomes to “manage risk to an acceptable level, based on the enterprise’s risk appetite with decision-making guided by a risk assumption model” [61]. A key inhibitor to making this happen is the lack of appropriate quantitative models that are familiar to business leaders as far as risk input from non-IT areas is concerned.

Other disciplines (specifically financial and actuarial) have mature risk assessment models. These areas need to be examined to determine their portability to IS risk management discipline; specifically to explore generation of quantitative information about risks associated with information. Of special interest should be the Value at Risk (VaR) method used to measure the risk of loss on a specific portfolio of financial assets [37]. Using VaR as a quantifiable measure of risk, the research hypothesis then becomes:

Given a system, the probability of a threat actually occurring (p_T), and a time horizon, VaR would be defined as a threshold value such that the probability that the loss associated with the system over the given time horizon exceeds this value is p_T .

One restriction in being able to pursue this stream of research is the absence of historical data that can be used to test this (and other) hypothesis. This problem is considerably alleviated by the control sensor based model presented in Chapter 6 [67].

References

- [1] R. Anderson, “Why Information Security is Hard- An Economic Perspective”, 17th Annual Computer Security Applications Conference, Dec. 2001.
- [2] Jonathan D. Andrews, “Erosion of Trust – E-commerce and the Loss of Privacy”, Information Systems Control Journal, Vol. 3, 2001, pp. 46-49.
- [3] Georges Ataya, “Risk-aware Decision Making for New IT Investment”, Information Systems Control Journal, Vol. 2, 2003, pp. 12-14.
- [4] Rudy Bakalov, “Risk Management Strategies for Offshore Applications and Systems Development”, Information Systems Control Journal, Vo. 5, 2004, pp. 36-38.
- [5] S. P. Bennett and M. P. Kailay, “An application of qualitative risk analysis to computer security for the commercial sector”, Eighth Annual IEEE ComputerSecurity Applications Conference, Nov.-4 Dec. 1992, pp.64–73.
- [6] Nicholas A. Benvenuto & David Brand, “Outsourcing: A Risk Management Perspective”, Information Systems Control Journal, Vol. 5, 2005, pp. 35-40.
- [7] B. Blakley, E. McDermott and D. Geer, “Information Security is Information Risk Management”, Proceedings of the 2001 workshop on New security paradigms, 2001, pp. 97-104.
- [8] Paul Brooke, “Risk Assessment Strategies”, Network Computing, 30th of October, 2000 (http://www.networkcomputing.com/1121/1121f32.html?ls=NCJS_1121bt).
- [9] S. A. Butler, “Security Attribute Evaluation Method: A Cost-Benefit Approach”, Proceedings of the 24th international

- conference on Software engineering, ACM, May 2002, pp. 232-240.
- [10] K. Campbell, L. A. Gordon, M. P. Loeb and L. Zhou, “The Economic Cost of Publicly Announced Information Security Breaches: Empirical Evidence from the Stock Market”, *Journal of Computer Security*, Vol. 11, 2003, pp. 431-448.
- [11] H. Cavusoglu, B. Mishra and S. Raghunthan, “The Effect of Internet Security Breach Announcements on Market Value of Breached Firms and Internet Security Developers”, *International Journal of Electronic Commerce*, Volume 9, Issue 1, 2004, pp. 70-104.
- [12] F. Cohen, “A Cost Analysis of Typical Computer Viruses and Defenses”, *Computers & Security*, Vol. 10, 1991, pp. 239-250.
- [13] Center for Internet Security (<http://www.cis.org>).
- [14] Drucker, Peter, 1988. ‘The Coming of New Organization’, *Harvard Business Review*.
- [15] Peter Drucker, “The Practice of Management”, Butterworth-Heinemann, 2007.
- [16] Criminal Take Control of CheckFree Web Site (<http://pcworld.about.com/od/networkin1/Criminals-Take-Control-of-Chec.htm>)
- [17] COBIT 4.1, ISACA (<http://www.isaca.org/>).
- [18] Enterprises Risk Management – Integrated Framework (<http://www.coso.org/>).
- [19] Covington, Prahlad Fogla, Zhiyuan Zhan, and Mustaque Ahamad. A context aware security architecture for emerging applications. In *Proceedings of 18th Annual Computer Security Applications Conference (ACSAC)*, pages 249–258, Las Vegas, NV, December 2002.

- [20] Risk Management: Implementation principles and Inventories for Risk Management/Risk Assessment methods and tools, European Network and Information Security Services.
- [21] G. Eschellbeck, "Active Security- A Proactive Approach for Computer Security Systems, Journal of Network and Computer Applications, No. 23, 2000, pp.109-130.
- [22] F. Farahmand, W. J. Malik, S. B. Navathe and P. H. Enslow, "Security Tailored to the Needs of Business", Proceeding of the ACM CCS BIZSEC, October 2003.
- [23] F. Farahmand, S. B. Navathe and P. H. Enslow, Electronic Commerce and Security – A Management Perspective, ISS/INFORMS Seventh Annual Conference on Information Systems and Technology, San Jose, 2002.
- [24] Shared Assessments, <http://www.sharedassessments.org>.
- [25] Todd Fitzgerald, "Ten Steps to Effective Web-Based Security Policy Development and Distribution", in Information Security Handbook, Harold Tipton and Mickey Krause Eds., Auerbach Publications, Boca Raton, FL, 2005.
- [26] Todd Fitzgerald, "Building Management Commitment Through Security Councils", Information Systems Security, May/June 2005.
- [27] D. E. Geer, "Making Choices to Show ROI", Secure Business Quarterly, Vol. 1, Issue 2, 2001, pp. 1-3.
- [28] A. K. Ghosh and T. M. Swaminatha, Software Security and Privacy Risks in Mobile E-Commerce, Communications of the ACM, Feb. 2001, Vol. 44, No. 2, pp. 51-57.
- [29] L. A. Gordon and M. P. Loeb, "Return on Information Security Investments", Strategic Finance, Nov. 2002.
- [30] John Hagerty, "The Governance, Risk Management, and Compliance Spending Report, 2008–2009: Inside the \$32B GRC Market", <http://www.amrresearch.com>.

- [31] M.H. Han, M. Y. Lee and T. H. Cho, “Fuzzy-Based Verification-Probability Determination Method for Dynamic Filtering in Sensor Networks”, IJCSNS International Journal of Computer Science and Network Security, VOL.8 No.8, August 2008.
- [32] Rebecca Herold, “Managing an Information Security and Privacy Awareness and Training Program”, Auerbach Publications, New York, 2005.
- [33] Board Briefing on IT Governance, 2nd Edition, <http://www.itgi.org>.
- [34] CISM Review Manual. 2006, Information Systems Audit and Control Association.
- [35] ISO/IEC 27005:2008, Information technology – Security techniques – Information security risk management, International Standards Organization.
- [36] ISO/IEC 27001, Information technology – Security techniques – Information security management systems – Requirements, International Standards Organization
- [37] Philippe Jorion, Value at Risk: The New Benchmark for Managing Financial Risk, 3rd ed. McGraw-Hill (2006).
- [38] K. Kark, Calculating the Cost of a Security Breach, Forrester Special Report, April 2007, <http://www.forrester.com>.
- [39] Soren Lauesen and Houman Younessi, “Six Styles of Usability Requirements”, Proceedings of REFSQ’98, Presses Universitaires de Namur, 1988.
- [40] D. Lilja, Measuring Computer Performance: A Practitioners Guide. 2000, Cambridge University Press.
- [41] U. Lindquist and E. Jonsson,, How to systematically classify computer security intrusions, IEEE Symposium on Security and Privacy, 1997, pp. 154 –163.
- [42] N. Mayer, A. Rifaut, and E. Dubois, “Towards a Risk-Based Security Requirements Engineering Framework”, 11th International Workshop on Requirements Engineering: Foundation

- for Software Quality (REFSQ'05), in conjunction with CAiSE'05, Porto, Portugal, June 2005.
- [43] N. Mayer, R. Matulevicius, P. Heymans, “Alignment of Misuse Cases with Security Risk Management”, Proceedings of the 2008 International Conference on Availability, Reliability and Security, pp. 1397-1404, 2008.
- [44] D. C. Montgomery, Design and Analysis of Experiments, 1991, Fifth Edition. Wiley.
- [45] R. H. Myers and D. C. Montgomery, Response Surface Methodology: process and product optimization using design experiments, 1995. John Wiley and Sons, New York.
- [46] Risk Management Components, Microsoft Press (<http://msdn.microsoft.com/en-us/library/cc500392.aspx>).
- [47] S.B. Navathe, G. P. Sharp and P. H. Enslow, “Assessing Damages of Information Security Incidents and Selecting Control Measures, a Case Study Approach”, Fourth Workshop on the Economics of Information Security, WEIS05, 2005, Kennedy School of Government, Harvard University.
- [48] Neumann, P. G., and Parker, D. B., A Summary of Computer Misuse Techniques. Proceedings of the 12th National Computer Security Conference, Oct. 1989, National Institute of Standards and Technology/National Computer Security Center, pp. 396-407.
- [49] An Introduction to Computer Security: The NIST Handbook, Special Publication 800-12.
- [50] OCTAVE Information Security Risk Evaluation (<http://www.cert.org/octave/>).
- [51] Applying OCTAVE: Practitioners Report, (<http://www.cert.org/octave/>).
- [52] OCTAVE Criteria, Version 2.0, (<http://www.cert.org/octave/>).
- [53] Octave Catalog of Practices, Version 2.0, (<http://www.cert.org/octave/>).

- [54] Information Asset Profiling, (<http://www.cert.org/octave/>).
- [55] Managing Information Security Risks: The OCTAVE Approach, (<http://www.cert.org/octave/>).
- [56] X. Parker and L. Graham, “Information Technology Audits”, CCH Inc., pp. 137 – 141.
- [57] Thomas R. Peltier, “Information Security Policies and Procedures: A Practitioners Reference”, 2nd Ed., Auerbach Publications, New York, 2004.
- [58] Thomas R. Peltier, “Information Security Risk Analysis”, 2nd Ed., Auerbach Publications, New York, 2005.
- [59] A. H. Phyto and S. M. Furnell, “A Detection-Oriented Classification of Insider IT Misuse”, 2004. Proceedings of Third Security Conference, Las Vegas, NV.
- [60] R. Plackett and J. Burman, “The Design of Optimum Multifactorial Experiments”, Biometrika, Vol. 33, Issue 4. June 1956. pp 305-325.
- [61] Mastering the Risk/Reward Equation: Optimizing Information Risks to Maximize Business Innovation Rewards, an industry initiative sponsored by RSA (http://www.rsa.com/innovation/docs/CISO_RPT_0808.pdf).
- [62] The SANS Policy Project, <http://www.SANS.org/resources/policies>.
- [63] H. J. Schummacher and S. A. Ghosh, “A Fundamental Framework for Network Security”, Journal of Network and Computer Applications, 1997, pp. 305- 322.
- [64] A. Singh and D. Lilja, “Criteria and Methodology for GRC Platform Selection”, To appear in ISACA (Information System Audit and Control Association) Journal. Volume 6, 2009.
- [65] A. Singh, A. and D. Lilja, “A Statistical Approach for Security Parameter Determination”, The 2009 International Conference on Security and Management.

- [66] A. Singh and D. Lilja, “Improving Risk Assessment Methodology: A Statistical Design of Experiments Approach”, 2009 ACM Conference on Security of Information and Networks.
- [67] A. Singh and D. Lilja, “STARTS: A Decision Support Architecture for Dynamic Security Configuration Management. 2009 IEEE International Conference on Engineering Management”, December 2009.
- [68] A. Singh, D. Lilja and G. Ray, “VIAP: A Framework for Managing Information Security Risks Due to Vendors”, Submitted to Information Security Journal.
- [69] Stephenson, Peter R., A Formal Model for Information Risk Analysis Using Colored Petri Nets, Technical Report of The Center for Regional and National Security, Eastern Michigan University.
- [70] Stoneburner, Goguen and Feringa. Risk Management Guide for Information Technology Systems, NIST Special Publication 800-30.
- [71] Sun Security Blueprints: <http://wikis.sun.com/display/BluePrints/Security+Blueprint>.
- [72] L. Trocine, and L. Malone, “Finding Important Independent Variables through Screening Designs: A Comparison of Methods”, Proceedings of 2000 Winter Simulation Conference, pp. 749–754.
- [73] Trusted Toolkit Policy Project, <http://trustedtoolkit.com/Documents/VendorThirdPartyPolicySample.pdf>.
- [74] K. Uehara, “J-SOX Challenge: Efforts to Comply With the New Japanese Regulation”, Information Systems Control Journal, Volume 5, 2008; pp. 34-37.
- [75] S. Viswanathan and R. Bhatnagar, “The Application of ABC Analysis in Production and Logistics: An Explanation of the

- Apparent Contradiction”, International Journal of Services and Operations Management, Vol. 1, 2005, No. 3, pp. 257-267.
- [76] Web Application Security Consortium, <http://www.webappsec.org>.
- [77] Charles C. Wood, “Information Security Roles and Responsibilities Made Easy”, Version 1, Pentasafe Security Technologies, Houston, TX, 2001.
- [78] J. Yi, D.Lilja, and D. Hawkins, “A Statistically-Rigorous Approach for Improving Simulation Methodology”, International Symposium on High-Performance Computer Architecture, 2003.
- [79] J. Yi, R. Sendag, L. Eeckhout, A. Joshi, D. Lilja, and L. K. John, "Evaluating Benchmark Subsetting Approaches" International Symposium on Workload Characterization, October 2006, pp 93—104.
- [80] Michael J. Cerullo and Verginia Cerullo, “Threat Assessment and Security Measures Justification for Advanced IT Networks”, Information Systems Control Journal, Vol. 1, 2005, pp. 35-43.
- [81] IRAM Risk Assessment Process, <https://www.securityforum.org/services/publictools/publiciram/>

Appendix A – Permissions to Reprint

- **Email granting permission to reprint figure Fig. 2.1.2.1 ISO 27005 Information Security Risk Management Process**

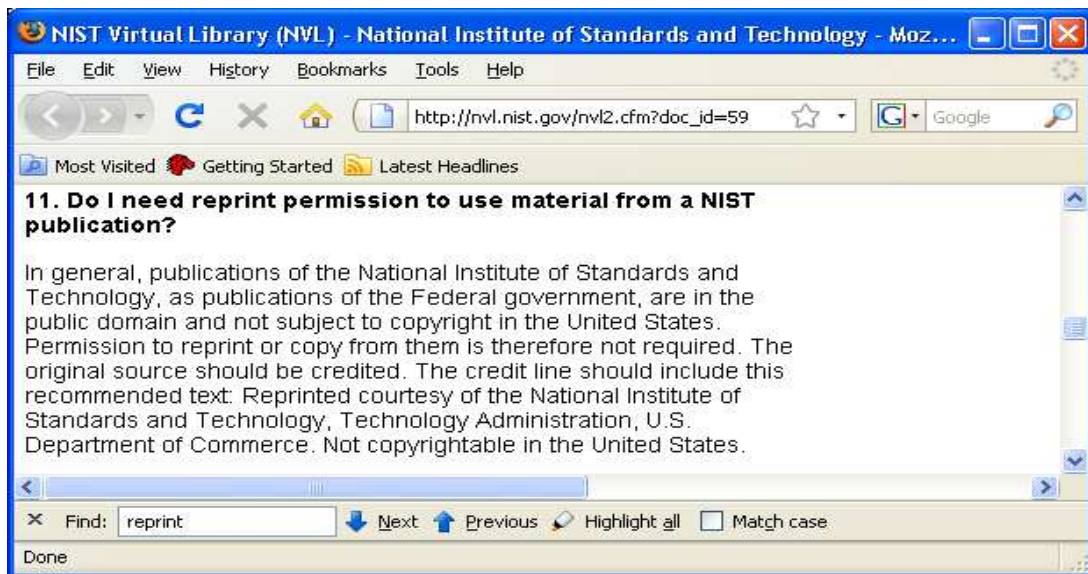
Dear Mr. Singh,
ANSI is granting you permission to use the material in your thesis. Include the following copyright statement:

This material is reproduced from ISO/IEC 27005 with permission of the American National Standards Institute (ANSI) on behalf of the International Organization for Standardization (ISO). No part of this material may be copied or reproduced in any form, electronic retrieval system or otherwise or made available on the Internet, a public network, by satellite or otherwise without the prior written consent of the ANSI. Copies of this standard may be purchased from the ANSI, 25 West 43rd Street, New York, NY 10036, (212) 642-4900, <http://webstore.ansi.org>"

ANSI wishes you success with your thesis and your future endeavors.

Regards,
Tim D [REDACTED]
Customer Service Manager
ANSI
Ph: [REDACTED]
Fx: [REDACTED]

- **NIST guidance on reprinting materials from its publications (applies to Figures 2.1.1.1, 2.1.1.2 and 2.8.1.**



Vita

Anand Singh was born in Varanasi, India. He was awarded National Talent Scholarship in 1990. He received his Bachelor's degree in Computer Science and Engineering from Indian Institute of Technology, Varanasi in 1996 and his Master's degree in Computer Science from Purdue University in 1998. He started his career as a senior technical architect responsible for performance of Cray SV1 platform. He has worked as a senior executive at Parametric Technology Corporation, US Bank and Target Corporation. He prides in his ability to bridge the gap between technology and business through strategic planning, development and implementation of cutting edge IT solutions for complex business problems. He is a highly sought after speaker on Information Security and IT issues. He is CISM (Certified Information Security Professional) and CISSP (Certified Information Systems Security Professional) certified.

Listing of Research Output

Conference Papers

- Anand Singh and David Lilja, “A Statistical Approach for Security Parameter Determination”, The 2009 International Conference on Security Management, July 2009.
- Anand Singh and David Lilja, “Improving Risk Assessment Methodology: A Statistical Design of Experiments Approach”, 2009 ACM Conference on Security of Information and Networks, October 2009.
- Anand Singh and David Lilja, “STARTS: A Decision Support Architecture for Dynamic Security Configuration Management. 2009 IEEE International Conference on Engineering Management”, December 2009.

Journal Papers

- Anand Singh and David Lilja, “Criteria and Methodology for GRC Platform Selection”, ISACA (Information System Audit and Control Association) Journal, Volume 9, 2009.
- Anand Singh, David Lilja and Gautam Ray, “VIAP: A Framework for Managing Information Security Risks Due to Vendors”, Submitted to Information Security Journal.

Invited Presentations

- “What are CSO’s Thinking About? Top Information Security Initiatives for 2008 and beyond ...”, Center for Education and Research in Information Assurance and Security (CERIAS), Purdue University, West Lafayette, IN, 01/30/2008, URL: http://www.cerias.purdue.edu/news_and_events/events/security_seminar/details.php?uid=ft0rhcapkbmtsod1lmsmlsam7o@google.com.

- “Information Security Focus Areas for Government”, 27th Annual Minnesota Government IT Symposium, 10/09/2008.