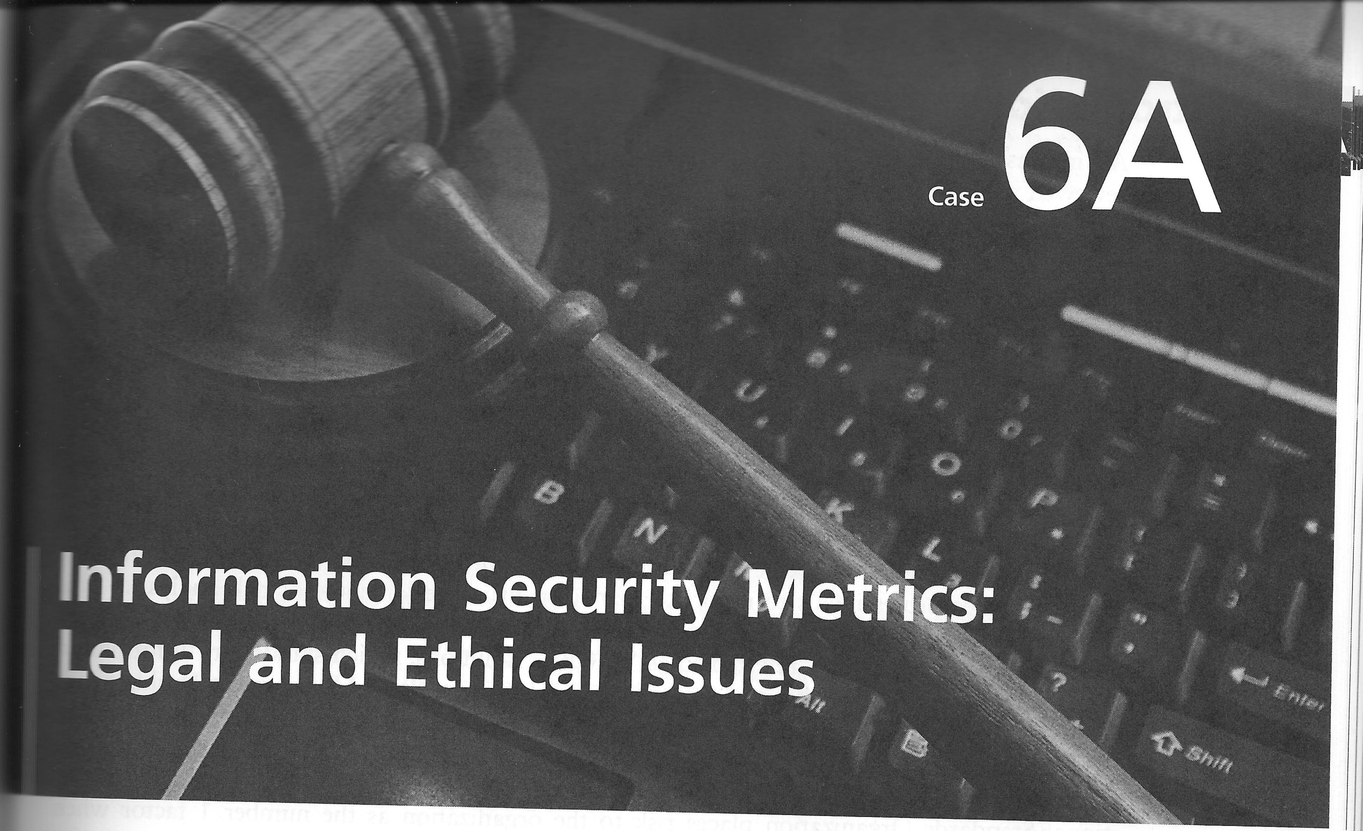

Case 6A

Information Security Metrics: Legal and Ethical Issues

Jennifer L. Bayuk
Stevens Institute of Technology

Jennifer Bayuk, CISM, CISSP, CISA, CGEIT, is an Information Technology due diligence consultant with direct experience in virtually every aspect of Information Security. Jennifer frequently publishes on IT Governance, Information Security Technology, and Audit topics, including two textbooks for the Information Systems Audit and Control Association. She lectures for organizations that include ISACA, NIST, and CSI. She is an industry professor at Stevens Institute of Technology and has master's degrees in Computer Science and Philosophy.

Overview

As security professionals move to the executive ranks, they have been advised to speak in business terms. The direct translation of security measures to business terms has to date been risk reduction.[1] So speaking in business terms has been to express information in terms of risk. It is very common to see panels of security experts nodding their heads in violent agreement that business risk is the most important factor to consider in designing security measures.

Introduction

An increasingly popular idea is that metrics with respect to information security should be reviewed at very high levels within an organization. There are publications that exhort members of boards of directors to ask questions about security and ensure that information security professionals are not censored as they present their information to influential decision-makers.[2]

The combination of these two notions, that the business language of information security is risk, and that information security metrics should be reviewed at high levels, has moved InfoSec into the realm of risk metrics. Risk management professionals are now called upon to use their skills to apply quantitative rigor to the InfoSec program.

This focus on risk has decoupled the honest presentation of progress in security measures from the metrics presented to executive management. Management reporting with respect to information security now focuses almost exclusively on risk to the organization rather than the goals and objectives of the information security program itself. Industry best practices exhort InfoSec professionals to report to management in the language of risk analysis.[3] This is evident given the plethora of products on the market that collect information security metrics that are collectively referred to as "ITGRC" solutions, which stands for "Information Technology Governance Risk and Compliance." They tout "Risk Management Dashboards" to be used for executive reporting, which abstract away details of program effectiveness and allow information security managers to arbitrarily color-code aggregate results so they can paint noncompliance of one target metric red while another appears green. Even the International Standards Organization places risk to the organization as the number 1 factor when considering security management: "This International Standard specifies the requirements for establishing, implementing, operating, monitoring, reviewing, maintaining and improving a documented Information Security Management System (ISMS) *within the context of the organization's overall business risks*."[4] Whether intentional or not, this risk focus often blurs facts concerning the true state of information security beyond recognition.

This paper describes metrics used to manage security in contrast with those often used to report on risk.

Security Metrics

Measurement is the process of mapping from the empirical world to the formal, relational world.[5] The measure that results characterizes an attribute of some object under scrutiny. Information security is not the object, nor a well-understood attribute. Attempts to create information security metrics fall into a wide variety of characterizations.[6] There are at least four types:

1. A—Activity-Related Metric: Metrics that Measure Work Activity
2. T—Target-Related Metric: Metrics that Have a Measurable Target (i.e., No Missing Logs)
3. R—Remediation Metric: Metrics that Show Progress toward a Goal
4. M—Monitor Related Metric: Metrics that Monitor Processes

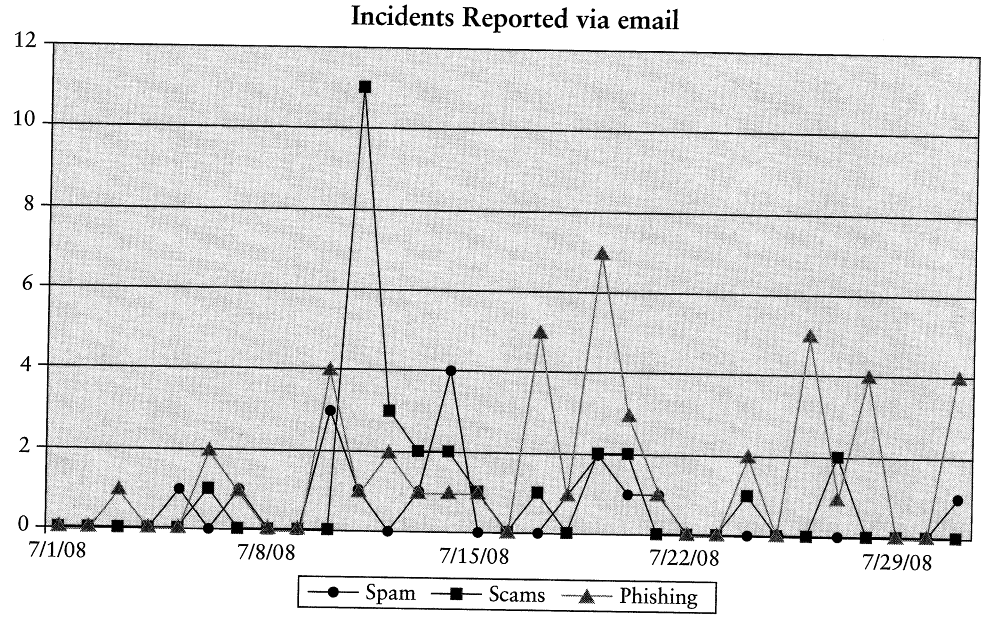

Note: Blank lines indicate no incidents were reported, mostly weekends.

Figure 6A-1 Activity Metrics

Courtesy Course Technology/Cengage Learning

Each type of characterization has its own utility in contributing information required to manage a security program.

Activity metrics are useful in resource allocation exercises. Figure 6A-1 shows activity metrics used to identify the number of incidents that security staff investigated over the course of the previous month. While there is no attempt at measuring the effectiveness of the incident investigation process, it is nevertheless useful in determining the number and type of activities that consume the time of security staff.

Target metrics are useful in verifying the correct implementation of technical security measures. The target is a total population that falls within the scope of a security measure. Figure 6A-2 provides an example of a target metrics where the target is a set of business unit applications and the attribute measured is whether they are all sending access logs to a central log collection server. Any deviation from 100% collected should immediately trigger an investigation. In this example, it appears that some technical difficulty was corrected midmonth, but resurfaced after a week or so of smooth operation.

Remediation metrics may be thought of as a subset of target where the scope of the target, thus the actual 100% measure, may not be known. Consider a case where an audit finding shows that user access is not properly terminated when employment is terminated because user ID strings are different in different systems and there is no correlation between those strings and any employee data record. A commonly agreed upon remediation for this situation is to deploy an identity management system to maintain such correlation. Figure 6A-3 is an example of remediation metrics that shows progress in the deployment of an identity management system. In this pre-production scenario, there are applications whose users are not yet

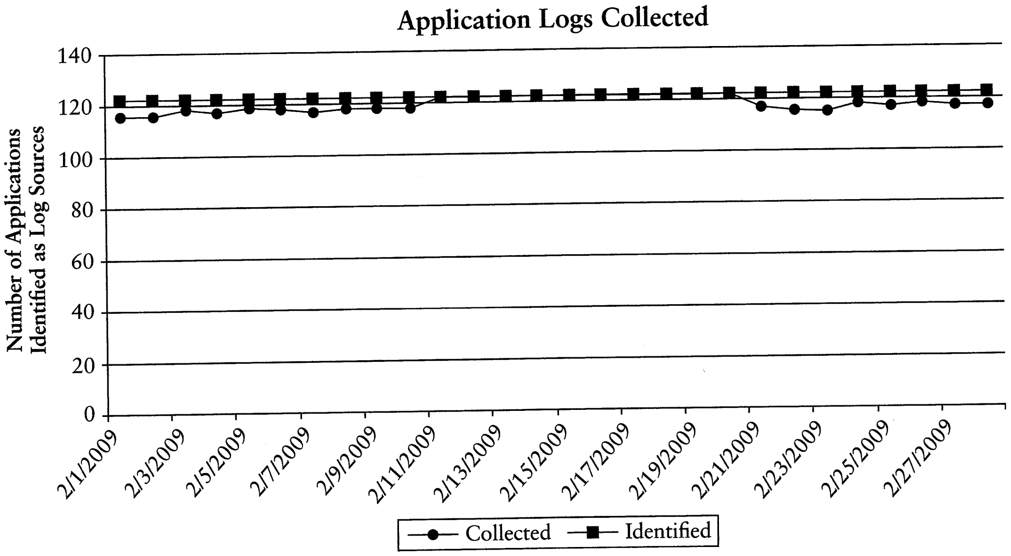

Figure 6A-2 Target Metrics

Courtesy Course Technology/Cengage Learning

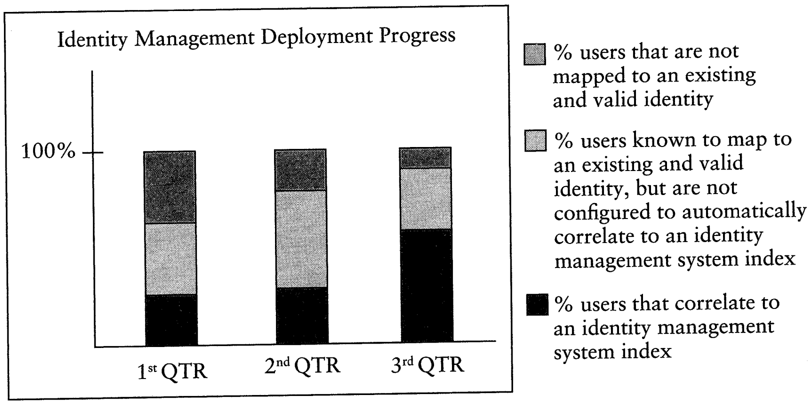

Figure 6A-3 Remediation Metrics

Courtesy Course Technology/Cengage Learning

identified by reference to an identity management system index, and they may or may not be mapped to an existing active and valid identity as the deployment continues. The number of users that are not correlated to the identity management system is considered a defect to be remediated and should become smaller as the deployment continues. A remediation metric is often used to determine the status of milestones in a security-related project plan.

Monitoring metrics are also useful in verifying the correct implementation of technical security measures, and in addition, they monitor security processes. Like target metrics, they have as a baseline the total population of a given universe. Figure 6A-4 illustrates how the correctness of firewall configuration may be measured. The top line represents the total number of firewalls in the organization. The line that sometimes dips below is the number of firewalls that were

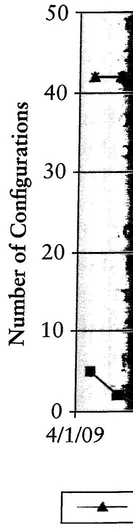

Figure 6A-4

Courtesy Course

available du
sent the nu
day. The nu
ifications th
verify both t
cuted. It ca
appears that
daily basis a

These basic
may be adap
in the realms

se Studies

Following an
onstrate to a
security pers
fact that ma
of the staff a
met. Auditor
ally met.[8] Th
three cases s
metrics.

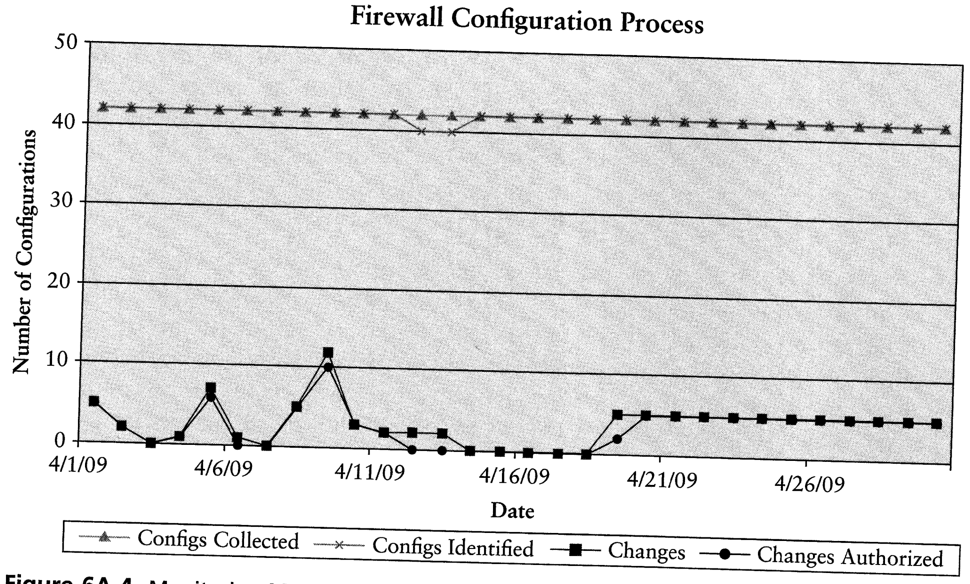

Figure 6A-4 Monitoring Metrics

Courtesy Course Technology/Cengage Learning

available during an automated daily configuration check. The lower lines on the chart represent the number of firewalls whose configuration check revealed changes from the previous day. The number of those changes is compared to the number of manual change control verifications that were performed by operations staff. A security manager can use this chart to verify both that automated controls are in place and that change procedures are properly executed. It can also be used to highlight systemic issues with a given process. In this case, it appears that there is a correlation between increases in the number of changes occurring on a daily basis and the occurrence of unauthorized changes.

These basic types of metrics are useful in managing a security program. Variations on them may be adapted to many situations and technologies encountered in day-to-day security tasks in the realms of policy, awareness and training, implementation, monitoring, and compliance.[7]

se Studies

Following are three case studies that show the types of metrics most used in practice to demonstrate to auditors and regulators that information security goals are met. The activity that security personnel engage in to ensure security goals are met is observed during audit. The fact that management monitors that process is followed is also important. But it is the *result* of the staff activity and management monitoring that provides evidence control objectives are met. Auditors must gather evidence that they can use to verify that control objectives are actually met.[8] That type of evidence comes in the form of remediation and target metrics. The three cases show alternative approaches to the process of gathering target and remediation metrics.

Case 1: The Question-and-Answer Approach

Metrics of this type rely on a network of risk management professionals to monitor progress in meeting some sweeping organizational objective such as compliance with regulatory reporting requirements. The risk management professionals themselves are not usually responsible for meeting the objectives, just for identifying the right set of individuals who should address them, organizing a reporting framework, and periodically polling those accountable to see whether they have achieved compliance.

These types of metrics are often created using online surveys. The major challenge in this approach is devising methods of identifying organizational structure and individuals within them that can be held accountable for survey answers. The state of the art in survey metrics collection utilities are systems wherein the accountable person can be linked to the control they are reporting on in such a way that when they log in to the risk application, their log-in is treated as a digital signature. This, combined with strict management policy that questionnaires must be answered within strict timelines, allow the risk managers to report the results with firmwide authority. Figure 6A-5 provides an example of the type of online surveys an accountable manager may encounter.

From such a survey, metrics such as Figure 6A-6 may be derived. Risk management professionals reduce accountability to self-reported task completion. From Figure 6A-6, it may be inferred that all of the new applications introduced into operation in the first quarter are using the firm's single sign-on system for authentication, and are compliant with security policy, design, and change control standards. However, there were a few introduced in the second quarter that did not make use of single sign-on and were not policy or standards compliant. The pie charts in Figure 6A-7 show this line of business (LOB) in comparison with the

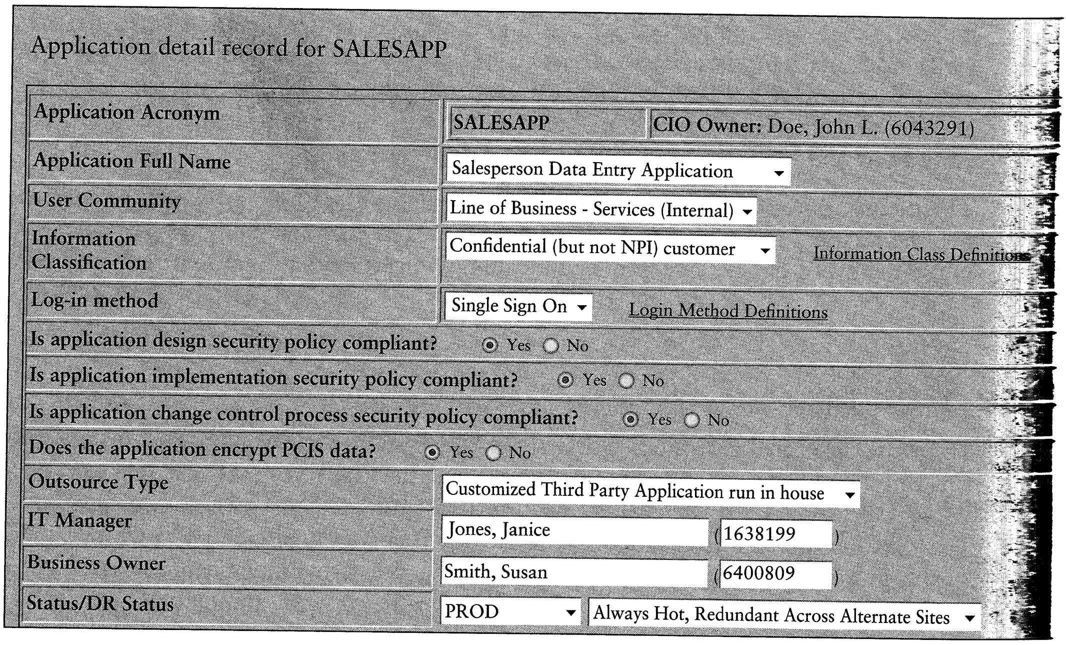

Application detail record for SALESAPP

Application Acronym	SALESAPP | CIO Owner: Doe, John L. (6043291)
Application Full Name	Salesperson Data Entry Application ▾
User Community	Line of Business - Services (Internal) ▾
Information Classification	Confidential (but not NPI) customer ▾ Information Class Definitions
Log-in method	Single Sign On ▾ Login Method Definitions
Is application design security policy compliant?	● Yes ○ No
Is application implementation security policy compliant?	● Yes ○ No
Is application change control process security policy compliant?	● Yes ○ No
Does the application encrypt PCIS data?	● Yes ○ No
Outsource Type	Customized Third Party Application run in house ▾
IT Manager	Jones, Janice (1638199)
Business Owner	Smith, Susan (6400809)
Status/DR Status	PROD ▾ Always Hot, Redundant Across Alternate Sites ▾

Figure 6A-5 Application Survey

Courtesy Course Technology/Cengage Learning

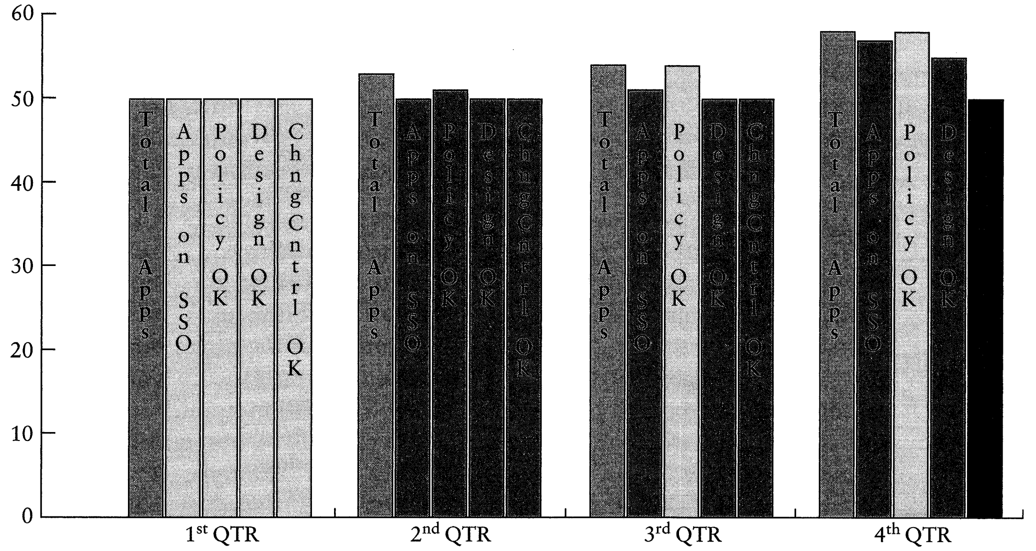

Figure 6A-6 Application Metrics

Courtesy Course Technology/Cengage Learning

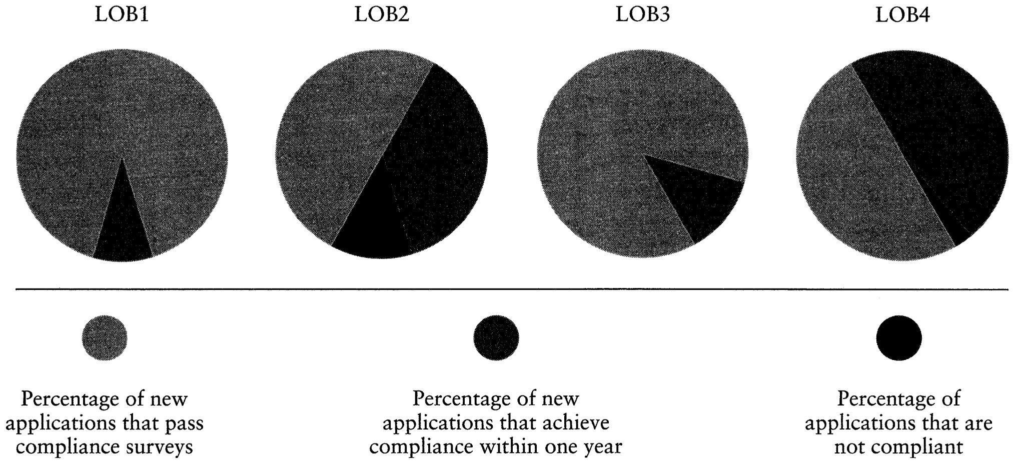

Figure 6A-7 Line of Business Application Metrics

Courtesy Course Technology/Cengage Learning

other LOBs in the firm. By comparison, LOB1 seems to have superior security deployment processes.

However, given that the data is based, not on automated measures, independent audit, or repeatable processes, but instead on data collection from accountable managers, all that can be inferred from Figure 6A-7 is that LOB1 *claims* to have fully compliant applications. The risk manager running the report is reduced to clerk status. These scenarios exist because the management objective on which the metric is modeled does not easily conform to available security measurement techniques. Suppose instead, there were automated verification procedures that could check to see if an application was in compliance, and there was a complete and accurate inventory of machines deployed in support of a given application. The inventory, in combination with automated verification procedures and/or technical configuration audits, could be used to establish that applications met firmwide security objectives. However, pressure for results in short timeframes sometimes require even the most technical of metrics professionals to resort to a question-and-answer method of data gathering.

These scenarios almost always suffer from real or perceived ambiguity in the way the "compliance" officer phrases overall goals and statements to be affirmed. For example, suppose an executive of a given organization is asked to affirm whether or not that organization has a policy to achieve compliance with a given regulation. The executive may sincerely respond in the affirmative, yet be thinking of a policy that states that the organization is fully committed to apply to all applicable regulation. At the same time, the executive may omit to mention that he or she does not believe that the regulation in question applies to that organization. Even when specific technology questions seem easy to verify, a respondent may find a question ambiguous enough for a misleading answer. For example, note the question in Figure 6A-5 that reads: "Does the application encrypt PCIS[9] data?" A manager who encrypts some but not all such data may honestly and misleadingly answer "yes." These examples and others like them leave a loophole of plausible deniability among those accountable for answering surveys while allowing the compliance officer to publish remediation metrics showing 100% compliance.

Case 2: The Risk Management Program

Risk assessment metrics with respect to security focus almost entirely on some adaptation of the risk management model prevalent in today's security literature:[10]

1. Assets
2. Threats to those assets
3. Vulnerabilities that may be exploited to enact threats
4. Probabilities that the exploits occur
5. Cost of controls that will reduce the likelihood the vulnerabilities can be exploited

Through a sound step-by-step approach, it is rare that an InfoSec professional will identify assets in any manner that would make sense from a business perspective. Performed correctly, the risk assessment would identify assets as products or services. However, there is overwhelming evidence in the security literature that assets are commonly construed to be computers themselves.[11] This makes threat identification easy, as a generic set of people who abuse computers can be easily identified: disgruntled employees, hackers, terrorists, and natural disasters. Vulnerabilities are also readily available as published sets of vulnerabilities from vendors as well as floods, fire, and hurricanes. Controls appear in this model as commercial off-the-shelf security products and services that will reduce the likelihood of threat.

So it is not too surprising that controls appearing as commercial off-the-shelf products and services are measured in similar ways by the InfoSec and risk management community. Automated tools can easily determine whether a given computer has a secure configuration. Figure 6A-8 demonstrates a metric commonly used in presentations of both InfoSec metrics and risk management metrics. It shows that there are a set of 27 computers in an organization all running the same operating system. There is an assumption that a secure

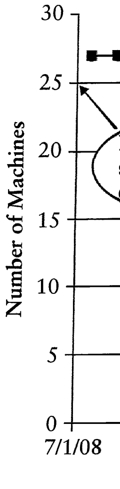

Figure 6A-

Courtesy Cours

configurati
puters in th
fall out of t
centage of
ating syste
shows a ta
support ch
graph is p
managemen

Figure 6A-
metrics and
place for so
organizatio
that one bu
business un
to explain
a risk reduc
that some I
cost of secu
probability
A security n

It is an ofte
reflect decisi
to determine
parameter, t
do a risk as

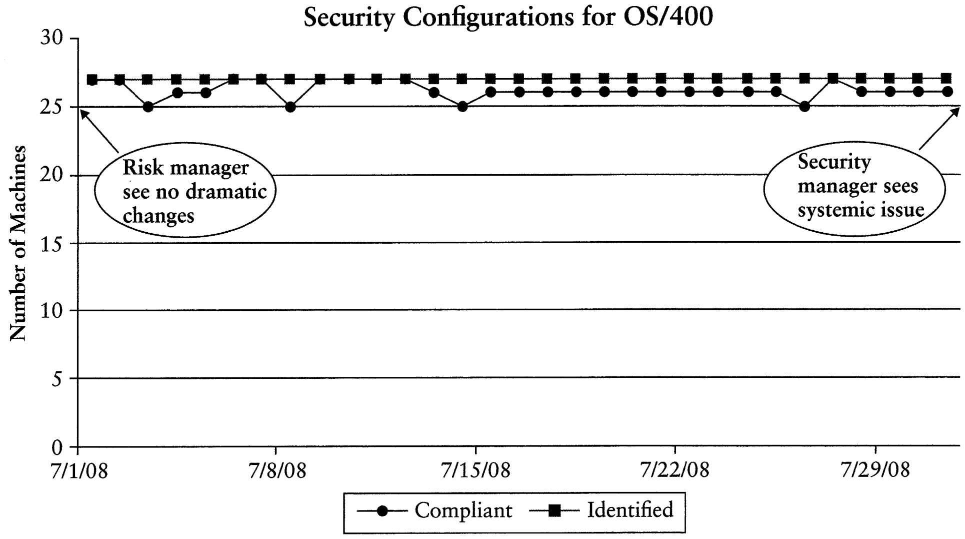

Figure 6A-8 Operating System Security Metric

Courtesy Course Technology/Cengage Learning

configuration of the operating system has been devised and is deployed on most of the computers in the set. However, it also shows that there are one or two machines that continually fall out of the secure configuration. To a risk manager, the lower line seems steady. The percentage of computers that are not secure seems low. The security configuration for the operating system as a whole seems stable. However, to the security manager, the same graph shows a target metric failed. The lower line indicates a weak link in the security operations support chain that could create a vulnerability to the entire platform. When this type of graph is presented to executives or auditors, it is almost always in the context of risk management.

Figure 6A-9 also demonstrates a metric commonly used in presentations of both InfoSec metrics and risk management metrics. It shows that some security control measure is in place for some type of platform that is deployed in different lines of business (LOB) in the organization. In both security management and risk metrics, it is intended to demonstrate that one business unit is more secure than another, to foster a sense of competition among business units. In the context of risk management, however, these metrics are instead used to explain different levels of risk and different approaches to the use of security controls as a risk reduction measure taken by different lines of business. Rather than walk away thinking that some LOB manager should be disciplined, executives are encouraged to compare the cost of security measures to the income produced by each line of business, and consider the probability that one of the known vulnerabilities indicated by the diagram may be exploited. A security manager seeing the same diagram sees a real and present danger.

It is an often glossed-over fact that the data presented in metrics like those in Figure 6A-8 reflect decisions based on judgments made by risk managers. An LOB may use Figure 6A-9 to determine that many of their noncompliance issues have to do with a single configuration parameter, that their machines are configured to "trust" others on their network. They may do a risk assessment on this configuration and decide that trust does not present a risk to

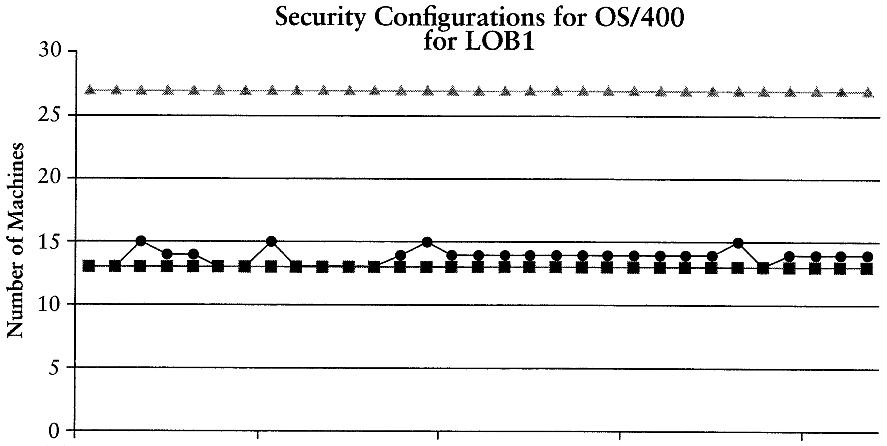

Figure 6A-9 Line of Business Platform Metric

Courtesy Course Technology/Cengage Learning

their environment. Where a security manager has been told that a LOB need not fix a vulnerability because their environment does not meet risk criteria necessary to fund a security fix, the group may mark that set of computers as the "equivalent" of having a compliant configuration. This may be recorded in a database that is consulted in the "automated" report that produces Figure 6A-8. To be clear on the potential impact of this risk management approach, the "Configured with Trust" line in Figure 6A-9 shows a number of computers that are not security policy compliant, but have been declared to be safe anyway due to a "risk acceptance" performed by line of business management. In Figure 6A-9, the 27 computers in Figure 6A-8 are represented as less than half policy compliant but half are not compliant due to a configuration that allows machines to trust each other. If a risk manager in the LOB finds the risk of trusting other computers acceptable in the course of their own risk assessment, the machines may be marked compliant. However, an auditor may only see Figure 6A-8 and not be told that different LOBs have different standards.

Moreover, even when configuration reports are produced with no exceptions, these risk management scenarios almost always concentrate on reported results of generic security software deployment packages. They rarely check to see if assets are protected through appropriate applications controls, or question whether management has allowed information to be disseminated too broadly rather than restricted by job function.

Case 3: The Sphere of Control

Organizations of this type have a centralized security function for most technology, while allowing some business units to independently manage business-unit-specific security measures. These distributed security functions usually have varying degrees of compliance with centrally published policies and procedures. However, when regulators and auditors come through, they are introduced to a person who is responsible only for security functions that overlap all business units, a *central* security officer, who proudly displays their metrics program while not mentioning that it does not cover 100% of the organization.

In contrast to the LOB-specific measures in Figure 6A-9, the central security function will gloss over technical differentiations in the data-gathering function. For example, measurements with respect to Windows operating systems are different than measurements concerning Linux. So normalizing these distinct measurements into industry-standard agreed-upon criteria for security measures seems like a commonsense exercise.

For example, Figure 6A-10 shows a set of security functions broken down by measureable features of operating systems that reflect them in a given environment. It does not necessarily matter that some operating systems implement security features better than others or that some instances of a given operating system are beyond the sphere of control of the central authority who generates the reports. It just demonstrates, whether, in the judgment of the central organization, these features are well implemented on the machines in scope. Such reports are meant to convey that the organization at least knows how to secure and measure things and that there should be no technology that requires security that would be beyond organizational expertise to deliver.

Central security organizations often feel justified in displaying this type of fuzzy target metric because they want to display true information but have been conditioned to allow the independent business units to evade their sphere of control. Where this is the case, they feel no responsibility to comment on the fact that their own management has allowed the independence. The lack of authority over LOBs could be perceived to indicate a lack of faith on the part of that management with respect to the security organization's ability to handle the special needs of the suborganization in the centralized program. The suborganization is acknowledged to be a stepchild under special care, like a child that is acknowledged to require special education whose efforts should not be averaged into the rest of the classroom because it would bring down the ratings for the school as a whole.

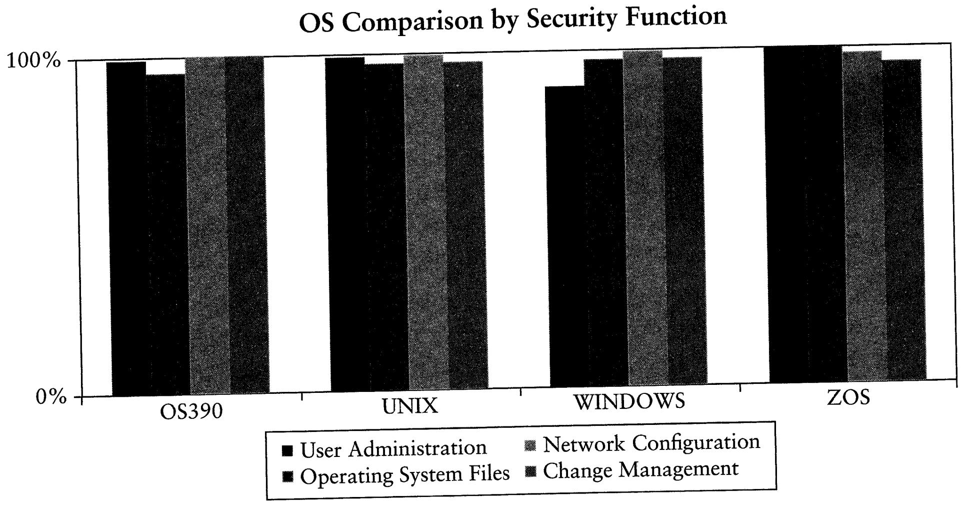

Figure 6A-10 Operating System Security Functions

Courtesy Course Technology/Cengage Learning

Conclusions

Let us suppose we are an InfoSec manager who needs to collect metrics to present to the board of directors (BOD). We are constrained on both time and level of detail. Suppose further we are new to the organization and not technical. We have no way to judge whether anything is secure or not for ourselves. In such a case, defining accountability and conducting surveys may be the only way to get the job done. The fact that the information is of dubious integrity is easily overlooked as the BOD meeting is short and is not expected to include technical details of how information was collected.

Even without the supposition that we are technical enough to evaluate information security risk ourselves, we are still constrained on both time and level of detail. Even when we know exactly where the problems are in the organization, it does not mean we have the solutions, especially when our sphere of control does not extend to LOB deployment programs. In this case, a security manager may feel that the program progress metrics are best served by presenting metrics as evidence of program evolution, that is, as charts that are clearly climbing uphill to meet targets. These normally leave the execs beaming about continuous improvements rather than upset about not meeting state-of-the-art, well-known, 10-year-old best practices.

A continuous improvement report where we do not meet targets is actually less preferable to present to the board than a presentation in which we are not at all expected to report on programs we do not control. Where we cannot even collect data from LOBs outside our control, we can simply omit mention of them and provide metrics on the deployment measures we fully control and endorse.

The metrics described in these cases will not tell the BOD whether the organization's information is secure. You know your information is secure when you have a security program and corresponding systems architecture. The way you know your program is working is via activity and monitoring metrics. The way you know your architecture is in place is via target metrics. When you find out you have a gap that leaves information exposed, you need remediation metrics. Metrics that correspond to a holistic security program demonstrate that management objectives for the program are (or are not) being met (or to what degree the gap). Without the explanatory guidance provided by the security program and information architecture, metrics represent collections of disconnected attributes, none of which are security. Every aspect of the security program need not be continuously measured to provide adequate assurance that the program is working. However, enough control points must be capable of being monitored often enough so that the combined measures of control points directly map onto the organization's requirements for information security. This mapping must be designed in such a way that there is capacity for demonstration that objectives are met.

Too often, target and remediation metrics absent program methodology and information architecture are mistaken for evidence of security program management. Executives on the receiving end of metrics often have a false sense of security. This type of blindspot, if it exists, does not extend to the rest of the organization nor even to its entire management team. Just as the product support team at Enron knew there was no broadband service offering,[12] the InfoSec managers and other technical managers know fully well that these types of metrics do not actually represent what they are implied to convey.

Endnotes

1. See
Sea

2. Boa
200

3. See
.secu

4. Ibid.

5. This
of m
Priva

6. See C
Auerl

7. For c
Thro
(ISAC

8. See th
publis

9. PCIS i
Card I

10. Such p
Critical
octave/
800-30

11. See Jaq

12. McLean
and Sca

Endnotes

1. See discussions of risk in Sherwood, John, Andrew Clark, and David Lynas, *Enterprise Security Architecture*, CMP Books, 2005.
2. *Board Briefing on IT Governance*, 2nd ed., IT Governance Institute (www.itgi.org), 2003.
3. See *Information Security Forum Standard of Good Practices*, SM3.3, www.securityforum.org. and BS ISO/IEC 27001:2005 BS 7799-2:2005.
4. BS ISO/IEC 27001:2005 BS 7799-2:2005.
5. This is a fairly generic definition of measurement, but for more reading on the concept of metrics as applied to security, see Herrman, Debra, *Complete Guide to Security and Privacy Metrics*. Auerbach, 2007.
6. See Cohen, Fred, *IT Security Governance Guidebook with Security Program Metrics*. Auerbach, 2008.
7. For comprehensive treatment of each of these realms, see Bayuk, Jennifer, *Stepping Through the InfoSec Program*, Information Systems Audit and Control Association (ISACA), 2007.
8. See the Certified Information System's Auditor's *Code of Professional Ethics*, a Standard published by Information Systems Audit and Control Association (www.isaca.org).
9. PCIS in this context refers to data required to be encrypted according to the Payment Card Industry Security Standards published at https://www.pcisecuritystandards.org.
10. Such publications are based on practices described by Carnegie Mellon in Operationally Critical Threat, Asset, and Vulnerability Evaluation$^{\text{SM}}$ (OCTAVE, http://www.cert.org/octave/), and National Institute of Standards and Technology in *Special Publication 800-30* (csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf).
11. See Jaquith, Andrew, *Security Metrics*, Addison-Wesley, 2007, pp. 235–237.
12. McLean, Bethany, and Peter Elkind, *The Smartest Guys in the Room: The Amazing Rise and Scandalous Fall of Enron*. Penguin Group, 2003.