



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

**The Controlled Event Framework for Information Asset
Security**

GSNA Gold Certification

Author: Chris Cronin, ccronin@chriscronin.net

Adviser: John C. A. Bambenek

Accepted: February 12, 2008

Table of Contents

1. Introduction.....	4
2. Purpose.....	5
3. The Problem: How Sensitive Information Poses Risks to Organizations.....	7
Document Handling Controls as a Risk.....	7
The Necessity of Risk.....	8
Example Risks	9
Current Document Handling Protocols	11
Legal Liability	12
The Three Essential Business Questions.....	13
4. The Framework.....	14
The Scope of the Framework	14
The Elements of the Framework	15
The Elements of the Framework: Events	20
The Elements of the Framework: Risk Categories	21
The Elements of the Framework: Instructions	23
5. Implementing the Framework	24
Being Reasonable	24

Organizational Commitment	25
Risk Assessment: Itemizing Risk Categories.....	27
Itemizing Events.....	30
Itemizing Formats.....	31
Developing Instructions.....	32
Incremental Adoption.....	35
6. Database Application.....	37
Why an Application?	37
The Tables.....	37
Time Based Security at the Event.....	40
Issued Instructions: Flexibility in the Framework.....	42
Issued Instructions: The Instruction Set.....	44
7. Auditing with the Framework	45
The Auditing Limits of the Controlled Event Framework	53
8. Conclusion	54
9. Glossary	58
10. References.....	61

1. Introduction

Organizations face multiple risks when handling documents that contain sensitive, non-public information. For example, personally identifying information (PII), proprietary secrets and other non-public information can be contained in a wide variety of information media and in any number of locations. Laws in several countries and 39 U.S. States require organizations to secure such information¹, and to properly inform individuals and organizations in the event that they expose that information, whether intentionally or unintentionally.

These laws often provide organizations the vague standard of “reasonable” to describe the lengths to which they must go to secure information, or to mitigate the risk to individuals whose PII was exposed.

But a number of questions arise when an organization considers how to meet this standard. How does an organization begin to control and monitor the security of non-public information? What is a reasonable control? Where is non-public information located in the organization? If organizations have security rules for document handling, are people obeying those rules? Is the organization even aware of the sensitive information in its possession? And what standards of security should be applied when documents are localized or distributed among jurisdictions with variable and multiple security and privacy rules?

While these and other security questions are asked by organizations, perhaps they are all covered by the one encompassing question: How can management and external authorities determine whether non-public information is well secured?

This paper proposes a framework for implementing, operating and testing document security controls within an organization. While much security management is meant to prevent people from doing things they ought not do, a framework is meant to help people do what they ought to do. In the case of the Controlled Event Framework for Information Asset Security, people are directed with some specificity on how to handle documents so they do their work effectively and securely.

Data Loss Prevention (DLP) is a focus of increasing concern to organizations that possess and handle non-public information. While DLP focuses on technical mechanisms for detecting and preventing the exposure of electronic information, the Controlled Event Framework provides for security through awareness.

Because security awareness is fundamentally important for sufficient security management, the primary security function of the framework is to formulate specific document handling instructions based on a risk assessment, then delivering instructions to document handlers to inform them of how to do their jobs securely.

This paper will present the case for a framework for document content security, and will present the proposed framework itself. Additionally, the fundamentals of a database application will be presented to help organizations understand the framework, and to recommend a mechanism for delivering it.

This framework will be useful to organizations that are trying to develop reasonable controls for reducing data loss by using a method for implementing and auditing controls.

2. Purpose

Countries around the world and an increasing number of U.S. states require by law that organizations that possess private information are also held accountable for the secure handling of that information. Because these laws are intended to protect individuals, they most often target personally identifying information (PII) such as Social Security numbers, account numbers for financial institutions, and medical information. However organizations are no less vulnerable to exposing other sensitive data, such as intellectual property, government secrets or corporate secrets, such as market moving information. Public reputation and tort liability surely pose risks to companies where regulations and laws do not. And for many organizations, such as defense contractors, national security is perhaps the most important reason for controlling documents securely.

Despite these liabilities, business and government organizations still are not doing enough to

implement Data Loss Prevention (DLP) controls². The scope of DLP is broad and the technologies that provide DLP controls provide varying methods for detecting and preventing the exposure of sensitive information while at rest or in motion. But DLP solutions do not work unless organizations carefully discover their requirements through a risk assessment and an analysis of how their documents and data are used. The Controlled Event Framework provides a comprehensive method for organizations to develop and implement these auditable requirements before they identify the tools they will use to assist them.

Controlling documents in a secure fashion is a difficult task for many organizations because of the many variables involved. The requirements for document control must extend to different technologies (paper, e-mail, data files, electronic documents), over multiple locations (file cabinets, brief cases, home computers, cell phones, servers, off-site storage facilities), among varying personnel (executives, their assistants, contractors, vendors), according to varying risk requirements (public, private, PII), and among multiple events (receiving, storing, printing, sending, copying, archiving, disposing).

Controls must not only state what must be done with documents during these events, but the controls must be communicated effectively to people who handle documents so they know what they can, cannot or must do with documents, and under what circumstances.

Finally, to determine whether or not an organization is securing documents well, there must be some way to test that these controls are effectively carried out. If organizations use a single framework for handling document security, these myriad issues can be well organized and documented, effectively communicated to involved parties, and tested by auditors.

The proposed Controlled Event Framework for Information Asset Security (hereinafter “Controlled Event Framework” or “the framework”) provides organizations with a structure for identifying needed controls and for communicating them to document handlers, administrators and auditors. The controls take the form of instructions that are provided to individuals. The instructions are developed by the organizations that adopt the framework and are based on meaningful security classifications, document types and document use.

The framework serves as a guide to adopting organizations by providing risk assessment tools, identifying needed controls, communicating controls, testing controls and organizing the entire control management process so that the DLP process is transparent and evident to internal or external auditing parties.

The framework does not provide specific controls that organizations should adopt. Controls and the security requirements that drive them must be derived from internal concerns. However, a data model, and illustrative security classifications and controls are presented to demonstrate how the framework would be executed within an adopting organization.

3. The Problem: How Sensitive Information Poses Risks to Organizations

Document Handling Controls as a Risk

While addressing the risks an organization faces in handling sensitive information, perhaps the most important thing to mention is what the author calls the “Arthur Andersen Rule.”

Put simply, the Arthur Andersen Rule posits that document handling policies are an added risk to organizations. Because document handling controls are meant to reduce risk, this rule seems contradictory. However, if there is no oversight, understanding, or consistent use of document handling policies, then the occasional compliance of - for example - destruction policies may appear to be capricious. A quick look at the Arthur Andersen case shows why this rule is so important to understand.

Arthur Andersen’s downfall began when its conflicted roles with their client, Enron, became public. But their fate was sealed when a jury found that the consultancy illegally destroyed documents that memorialized their conflicted work with Enron³. Their verdict was based on an appearance that Arthur Andersen “corruptly” gave orders to destroy documents that would serve as evidence against them. While Arthur Andersen’s internal document retention policies allowed them to destroy documents under certain circumstances those documents were among the few that were ever subjected

to the destruction policy. As well, Arthur Andersen's document destruction policies prohibited them from destroying documents that related to an engagement which had a law suit pending. Arthur Andersen's management defended themselves by claiming ignorance to this important caveat. The jury found the document destruction to be capricious and corrupt, and therefore an obstruction of justice.

Had Arthur Andersen not had a document retention and destruction policy, management may never have thought to destroy their Enron documents at all. Had document handlers known the important caveat that disallowed the destruction of documents, their damage would have been reduced. Hence the Arthur Andersen rule: Document handling policies are an added risk.

The Controlled Event Framework is meant to reduce risks to organizations that adopt it, so the framework has been designed with oversight (in the form of internal audit) as an essential attribute. The details of this will become apparent as the framework is presented. Without an organizational commitment to oversight and the resources that are required for oversight, the framework would most likely fail.

The Necessity of Risk

If risk is defined as the combination of threats, vulnerabilities and impact, then any organization that handles sensitive information must accept some risk. Without possessing and using some form of sensitive information, most organizations would not be able to function. Whether that sensitive information is the language contained in confidential contracts, employment agreements, compensation, membership lists, client data, intellectual property, just about any conceivable organization has information that poses a liability.

While this point may seem obvious, it is essential to remember when considering the degree to which an organization will protect information. Consider, for instance, a university. Among the several kinds of documents a research university will have in its possession are reputational data (student grades), PII (student loan account numbers), public information (course curricula), intellectual property (research and development) and occasionally government secrets (a subset of research and

development). When a university decides that it must control the documents in its possession, it will likely decide to not apply the same level of security and scrutiny to all its documents. The cost and inconvenience to the university to secure its curricular material to the degree that it must secure its government-funded research and development would be prohibitive.

So part of an organization's risk management is to classify the documents and data in its possession by the type of liability they pose, and to do so in a way that is meaningful to the organization. This will be discussed more thoroughly in the Risk Assessment subsection of Section 5, "Implementing the Framework."

Example Risks

For the purposes of the Controlled Event Framework, we define risk as the combination of three elements: a threat, a vulnerability and an impact. Risk is numerically calculated in various ways, depending on the risk management method. But generally, a risk is low or high depending on how low or high its elements are. For instance, as in the above example, the impact of a hacker copying course curriculum files from a publicly accessible server is low, so the risk is low. However, if the names of faculty and students who are working on a sensitive research project for the Pentagon are published on that server, the risk is higher.

As the framework and its implementation are described in this paper, example risks will be used to demonstrate some of the framework's attributes. The elements of these risks will be:

Threats: Events that cause organizations to lose control of information or information systems.

- Inconsistently used document handling policies (a threat and a vulnerability)
- Theft or loss of end-user data devices (laptop, backup tape, PDA)
- Carelessness with sensitive information (PII on web pages, printed copies left behind in a public space, accidental e-mail attachment)
- E-Discovery (response to subpoenas)

- Hacking/espionage/sabotage
- Malicious exposure from inside (selling competitive data, revenge by a disgruntled employee)

Vulnerabilities: Attributes to an organization or systems that allow threats to succeed.

- Inconsistently used document handling policies (a threat and a vulnerability)
- Lack of oversight for compliance with handling policies
- Lax document handling policies and procedures
- Insufficient systems security
- Staff who are unaware of security risks or handling policies
- Unencrypted data devices
- Ignorance of what sensitive documents or data are in possession
- Ignorance of when exposures of sensitive information occur
- Ignorance of legal requirements for responding to exposed sensitive information
- No incident response procedures
- Incident response procedures that are not followed
- Inadequate technical controls
- Engaged third parties who do not have adequate document handling controls

Impacts: The negative effect on people or organizations when a threat succeeds.

- Exposing individuals to fraud, security breaches, or harm to reputation
- Exposing organizations to risk (loss of competitive advantage, regulatory violations, security vulnerability and breaches, exposure of intellectual property)

- Violation of regional laws: Currently, 39 US States EU member nations, Australia and Canada all have laws that obligate organizations to handle nonpublic data with heightened care
- Violation of contractual obligations
- Violation of regulations and standards (HIPAA, Gramm Leach Bliley, Sarbanes Oxley, PCI DSS etc.)
- Civil action, either by clients or affected individuals.
- Injury to public reputation.

These lists are only illustrative and not intended to be comprehensive. For instance, organizations that are concerned about regulatory compliance while handling non-public documents may have more than one set of regulatory standards to be concerned about.

We will discuss how these lists are used in the implementation phase.

Current Document Handling Protocols

The most challenging part of an organization's effort in establishing a DLP program is in finding a model that works for them⁴. Perhaps the best known classification method for information is the method used by governments. In the case of the United States, the terms, *Top Secret*, *Classified*, and *Declassified* are easily recognizable, but perhaps not useful to most organizations. For good reason an employee may feel foolish for recommending such stylized classification names, because while they sound familiar, what do they mean? And what about the more generic classification methods that use *High*, *Medium* and *Low* or even *Public*, *Private*, *Sensitive*? Even if these classifications seem more benign than the dramatic government classifications, how are they meaningful to an organization that uses them?

Other information categorization methods direct practitioners to categorize documents by specific information types. Categories called *financial reports*, *vendor information*, *contracts* could be

found in many businesses. But once those information category types are established, a seemingly endless list of new document types including *employee information in a database*, *employee information in a file drawer*, *R&D e-mail* would grow as conscientious staff members consider other likely risks they habitually face. Now imagine the myriad rules and administrative time such a categorization would likely require. The risk that employees will not adhere to security rules increases as the complexity of the rules increases.

The State of California provides guidance to Information Systems Employees in a document titled, “California Counties ‘Best Practices’ Information Security Program.”⁵ Their guidance for creating security categories is to create a classification that is meaningful to the organization. Security categories can be created by an organization when they ask; what is the risk associated with the document, and can information handlers easily determine categories based on some quality of the information.

The Controlled Event Framework uses a similar approach to this suggested method as it provides organizations the best chance of producing categories and rules that are meaningful to its members.

Legal Liability

While legal liability is implied in the above list of impacts, it deserves special consideration. Privacy laws vaguely require organizations that expose data to respond in “reasonable” ways, or to apply “reasonable” security measures. However, organizations who possess PII may not know what “reasonable” means. Nor may they agree with external authorities about the meaning of the word when defending themselves after an information breach occurs.

The framework addresses the issue of “reasonableness” in two ways: one, by providing an organization with a way to determine reasonableness based on risk; and two, by demonstrating an effort to apply the best controls that the adopting organization could afford in time and money.

If organizations are to be held legally accountable for properly securing information, or for responding to exposures of that information, it must know whether their controls are reasonable, and be

able to demonstrate how they know. The framework satisfies both of these needs.

We will discuss later in the paper the role that laws and regulatory compliance play in developing the framework, but we should briefly note here that as an organization adopts and implements the Controlled Event Framework, they will need to know what government agencies require of them while securing information.

The Three Essential Business Questions

Despite the myriad variables and complexities involved in data loss prevention, an organization needs to understand its information risks and to know that the risks are being managed appropriately.

This paper poses three essential business questions that every business leader must be able to answer in order to know that their information security risks are appropriately managed:

1. Are our security controls reasonable? The purpose of the question is to balance effectiveness of controls with the cost of controls. “Reasonable” is a term that laws and courts use ambiguously to determine whether security was what it should have been when culpability in an exploit is being determined.
2. Are our security controls understood? This question reminds managers that the individuals who use documents or data must know the rules they must abide by, or at least have those rules at hand. The framework will provide a high degree of security awareness by providing detailed instructions that fit specific scenarios.
3. Are our security controls working? To prevent an avoidable exposure or the “Arthur Andersen Rule” an organization needs to know whether the rules are functioning the way they were designed to.

As the framework and its implementation are described, these three questions will be revisited to demonstrate the bottom-line benefit that the framework presents to organizations that adopt it.

4. The Framework

The Scope of the Framework

The Controlled Event Framework is designed to increase information security by communicating clear and appropriate handling instructions to the right people at the right time. The framework does not make assumptions about what controls or risks are appropriate for the organizations that adopt it, nor does it provide specific controls. The reason for the non-specificity of the framework is to provide organizations with a process for control, not the means for control. By necessity, the means an organization chooses to classify and control its information must be driven by its risks and capabilities.

The framework also prepares organizations who wish to implement DLP systems to help them evaluate the products that are available to them before they commit to a solution. When DLP systems providers and experts discuss securing documents and information, they focus early on classifications of information.

Classification schemes are a categorization of what information an organization carries, and what risk is associated with each category. Choosing a classification scheme is an essential early step in establishing document security, but it is essential that the risk categories are meaningful to the organization.

Organizations possess a wide variety of information types. Retail chains, for instance, can rely on well-formed credit card numbers and checking account numbers to enter, pass through, and reside in its point-of-sale systems as data. The chain's employee PII will likely be in predictable locations and in predictable formats. These retailers will have well-formed, inconsistently formatted price lists from vendors as data or documents. Loosely formed and inconsistently formatted terms and conditions will reside in contracts, and equally variable strategy documents will proliferate through systems, briefcases and homes.

Given these varieties of information and document types, a retailer can still with some ease categorize its documents based on the risk to customers, vendors, employees and the business. After

Chris Cronin

these risk categories have been identified, the retail chain reduces its risk of discovering an unknown, uncontrolled type of nonpublic information in its possession.

This is common in many organizations. The lifecycle of information can be fairly predictable for a large amount of the information that some organizations use. Correctly classifying and handling the documents in such situations can be relatively straightforward.

But as the framework is described in this paper, we will consider the work of a professional services organization that must handle documents from many types of organizations. Professional services organizations, such as law firms and consultancies, have very complex document control requirements because they often gather documents from clients who work in a variety of businesses.

At any given moment in a law firm, highly sensitive PII, reputational information, intellectual property, business strategies, market moving data, and even government and defense secrets can be resident on paper, electronic documents, databases, e-mail, or audio and video media. This same variability should be expected in consulting firms.

But unlike the retail chain example above, these professional services firms cannot easily predict what non-public information they will obtain at any given point, where it will reside, or in what form. This variety of obtained information also reduces the likelihood that the firms understand the obligations that each data type mandates, whether the mandate is legal or otherwise.

The Controlled Event Framework will be described using examples primarily from professional services firms to demonstrate its flexibility, and to demonstrate why a carefully paced adoption of the framework may be necessary and legally sound.

The Elements of the Framework

The Controlled Event Framework is comprised of four elements: events, risk categories, formats, and instructions.

Events: Events describe things that people or systems do to documents or data (including leaving them alone). Controls, and their related instructions, are attached to events by telling people how a document or data should be handled while executing an event.

Risk Categories: Risk categories are classifications that imply a level of security rigor that must be applied while securing information. For instance, a large manufacturer may apply more scrutiny to securing the Social Security numbers of its employees than to its R&D documents. That is, if it is more concerned with the legal requirements of securing PII than its commercial needs to secure its intellectual property. In such a case the manufacturer may have at least two categories: PII and Intellectual Property.

Formats: Formats describe the technical or physical “container” of information. A document image that contains a list of IP addresses for Pentagon networks is significantly different from a database file that contains the same information because the rules for handling the documents will be different. Recall Essential Business Question Number 2, “Are our security controls understood?” Instructions for copying data from databases would describe queries, reports and exports that are either disallowed or would have prescribed methods. However, such rules would be nonsensical, or variably interpreted by a person who was assigned to extract data from a PDF file.

Instructions: Explicit rules that direct document users in how to handle documents during an event. These instructions are also provided to administrators and auditors. The instructions can provide the detail of handling rules, or refer to existing policies and procedures. Instructions increase information security by guiding people toward the predictable actions that have been designed to reduce risk.

We will also see how the development of instructions is directly tied to countering known risks, such as stolen laptops, unauthorized use and hacking.

While the framework has these four elements - events, risk categories, formats and instructions - it can be quickly understood when represented in the form of a lifecycle. While documents generally may not have such a well-regulated and consistent life as the framework implies, events do tend to tell an organization what to expect of a document’s use from the time it is received or created until the time it is disposed of.

In Figure 1, the Controlled Event Framework is represented in a generic form, showing how risk categories can drive levels of security rigor in handling instructions.

Figure 1

	Event 1	Event 2	Event 3	Event 4
Standard Security Category	Standard handling instructions and oversight	Standard handling instructions and oversight	Standard handling instructions and oversight	Standard handling instructions and oversight
Heightened Security Category	Heightened instructions and oversight	Heightened instructions and oversight	Heightened instructions and oversight	Heightened instructions and oversight
Strict Security Category	Strict instructions and oversight	Strict instructions and oversight	Strict instructions and oversight	Strict instructions and oversight

This generic display of the framework presents an abbreviated list of events (many organizations will require many more, as we will see below), three levels of security requirements, and instructions associated with each event.

An event, such as *Copy* or *Dispose* will have instructions for copying or disposing documents or data that will support the level of risk associated with them.

In Figure 2, we again see an abbreviated set of events to demonstrate another point. People who handle documents, administrate the support of document handling, or audit compliance with document handling security, will need to know how documents of various formats should be handled, given a risk category.

Figure 2

	Categorize	Receive	Copy	Dispose
Paper	Follow "Assign Risk Category" procedure	Inventory in document log and place in indexed banker box	Scan according to "Scanning Documents" procedure.	Return paper to client. Run "Certified e-Shred" procedure
Electronic document image	Follow "Assign Risk Category" procedure	Inventory in document log and place in controlled file directory	Do not copy	Run "Certified e-Shred" procedure
Electronic document	Follow "Assign Risk Category" procedure	Inventory in document log and place in controlled file directory	Do not copy	Run "Certified e-Shred" procedure

Figure 2 gives a representation of instructions that an individual or team would see when they are assigned responsibility for documents in a given risk category. Figure 2 is an important representation to keep in mind. During the course of explaining the framework and its components, many rules, events, document types and risk categories will be described, and the count of instructions that they require will be high. But while implemented, the Controlled Event Framework provides people who handle documents with a customized, short list of instructions that is distilled for their purposes.

And finally, as a three-dimensional representation, the framework can show how event handling rules vary by format and risk category.

Figure 3

	Strict Security	Categorize	Receive	Copy	Dispose	
		Follow	Inventory in	Scan	Return paper	
	Heightened Security	Categorize	Receive	Copy	Dispose	st. Run "Certified e-Shred" procedure
		Follow	Inventory in	Scan	Return paper	ertified "red" edure
	Standard Security	Categorize	Receive	Copy	Dispose	ertified "red" edure
		Follow	Inventory in	Scan	Return paper	ertified "red" edure
Paper		Follow "Assign Risk Category" procedure	Inventory in document log and place in indexed banker box	Scan according to "Scanning Documents" procedure.	Return paper to client. Run "Certified e-Shred" procedure	ertified "red" edure
Electronic document image		Follow "Assign Risk Category" procedure	Inventory in document log and place in controlled file directory	Do not copy	Run "Certified e-Shred" procedure	ertified "red" edure
Electronic document		Follow "Assign Risk Category" procedure	Inventory in document log and place in controlled file directory	Do not copy	Run "Certified e-Shred" procedure	

As a set, instructions for handling documents of different formats would be created within an assigned risk category. For each risk category in an organization, the handling instructions for a document format would reflect the security rigor associated with that risk category.

In other words, paper documents will have a set of *Copy* instructions for each risk category. These copy instructions will permit, prescribe, or deny actions based on that risk category.

As we will see later in the paper, these instructions are not meant to imply uniformity in all handling instructions, but to state the weakest security rigor allowed for a document format during any event within any risk category.

If the purpose of the framework is to help people act in a secure way while doing their work, then applicable instructions are essential to the success of security controls. These instructions should be;

- a) Tailored to the rigor of security that a document's information exposes an organization to

- b) Applicable to the document format
- c) Bound to an event
- d) Should provide the least level of security rigor that the organization requires of the event

The Elements of the Framework: Events

The list of events that can occur to documents and data will be largely the same from one organization to another. Except for special cases, most organizations are likely to do the following to documents during their lifetime:

1. Categorize information by risk
2. Receive a document
3. Store a document
4. Copy
5. Create (either new information or a new document)
6. Use (Read, edit, analyze, etc.)
7. Convert (Includes print, export, scan, OCR)
8. Backup
9. Send
10. Dispose
11. Archive

Specialized handling events that may be important in organizations include: vetting or editing documents for sensitive information, cleaning documents to remove meta-data, serializing documents (such as Bates stamping), watermarking documents, etc.

When an organization implements the controlled event framework, it will identify the handling events that are significant to it. The staff members who implement the framework must ask what

information handling events the organization will want to control.

The Elements of the Framework: Risk Categories

A risk category is an element of the controlled event framework that delineates a level of security rigor within which events must be controlled. Perhaps the best known type of risk category for information is government document classification systems; e.g. *Classified*, *Secret*, and *Top Secret*. This model of document classification is often used when consultants or security writers offer suggestions for securing documents. The model considers security by using two parameters: who is qualified to access the information; and what are the access rules for a given document. It is a classification scheme that is well-suited to hierarchical organizations where information access is associated with rank.

However, organizations should use methods of information classification that are aligned with their specific risk issues. For this reason the Controlled Event Framework uses the directing language “Risk Category” and not the vague word “classification” to describe the level of security that is associated with documents and their handling instructions.

By associating handling instructions with an organization’s risk, the organization can align its administrative costs with the risks it is trying to mitigate. This process will become an essential part of how an organization can say its information security controls are “reasonable.” We will explore this later as we look at implementing the framework.

Let us consider for a moment the business reason for classifying information according to risk. The example given below in Figure 4 provides a slice of the controlled event framework in a law firm. The one event presented is *Copy*. The law firm in this example uses three risk categories: *Standard Security*, *Confidential* and *Personal Information*.

Figure 4

	Copy
Standard Security	Log copies
Confidential	Only Pre-approved copies
Personal Information	Prevent copies

Now consider the administrative costs, in time and money that the law firm must absorb while handling documents according to this event and these risk categories. For the purposes of the example, we will make the unlikely assumption that only electronic documents are being controlled, and that all electronic documents are resident on a system that logs handling events automatically.

The law firm determined that it had three risks related to the documents it handles:

- What will happen to our reputation and business costs if standard business documents (employee policies, public documents, internal administrative documents) are exposed due to careless handling or hacking?
- What will happen to our reputation, business costs or our clients if their confidential business information in our care is exposed?
- What will happen to our reputation, business costs or our ability to do business if personally identifying or reputational information is exposed?

This law firm preliminarily decides to consider three risk categories according to the business risks it identified.

Now imagine that the law firm works mainly with publicly available documents, such as real estate transactions. And a division of the firm works on corporate mergers. Perhaps two partners in the

firm work on medical class actions.

If all of the documents in the firm's care are handled by the *Standard Security* instructions, then any number of copies of documents is allowed to be created. If a dozen copies of a land deed are leaked, there is little damage to the firm or the client - except for some embarrassment – because land deeds are public documents. But this law firm also handles medical class actions. If copies of patient testimonies - containing PII - are copied and sent to an attorney's laptop then they are harder to track down when the firm disposes of those documents after the case closes. Even if all copies of the PII documents were destroyed at the server, the firm's risk of exposing the documents from a laptop theft or accidental e-mail attachment remains.

But now imagine that the firm decided to handle all documents at the highest risk category to be sure that all documents are handled with the highest level of care. If the firm's copies of land deeds are subjected to the same high security rigor as PII data sets, then the firm must never copy a land deed. Further, they must invest in oversight that demonstrates that they never copy a land deed. Carried over to all events, the law firm would have to invest a lot of time and money to mitigate low risk.

Risk categories are important because they help an organization, and outside agencies that inquire, to determine whether handling instructions for documents are reasonable, i.e. that they were aligned with risk.

The Elements of the Framework: Instructions

Finally, the framework provides instructions to users, administrators and auditors for handling documents according to security rules. As representations of the previous elements show, instructions are provided for each combination of an event and format, and adapted to a risk category. Essential to all instructions are three aspects: *who*, *what* and *where*.

Instructions must state who may (or may not) do what and where. For example, a good printing instruction at a university may read “Only bursars may print detailed student loan reports and must do so at printers in the locked file room.” A good disposal instruction at an engineering firm may read “The lead engineer or the lead engineer's designee must shred all super-ceded draft designs using the e-

doc file shredder and store all shred reports in the design plan log.”

It is also a good idea to reference within instructions the policy or procedure that provides fuller, more detailed language as to how these instructions must be performed.

The purpose of instructions is to provide awareness to people who handle documents, who administer systems support, or who audit the appropriate handling of documents. If an organization’s members know who can do what and where while handling a document, then ambiguity has been reduced, awareness has increased and security has increased.

During the “Implementing the Framework” section of the paper, we will see how distinct instructions will be applied for users, administrators and auditors, and how these instructions will be designed, recorded and delivered.

5. Implementing the Framework

Being Reasonable

In Section 4, the paper discussed legal liability in terms of “reasonableness.” In an article entitled, “Notification of Data Security Breaches” Paul M. Schwarz and Edward J. Janger demonstrate that the word “reasonable” is interpreted by the courts giving organizations under scrutiny much leeway in defining for themselves what “reasonable” measures are. However, they also describe a problem that outside authorities have in determining how an organization came to their understanding of what “reasonable” meant to them, and how they determined that their controls were reasonable.⁶

Certainly an organization that has not made efforts to control non-public information will fail to prove to anyone that they have taken reasonable measures to secure information, or to respond to unauthorized exposures of information. However, an organization that is implementing control procedures, measuring their effectiveness and improving controls, will very likely fit inside that realm of leeway that Schwartz and Janger describe. The implication is that simply demonstrating effort to

apply and improve security controls around non-public information mitigates liability better than no attempt at all. Organizations that feel overwhelmed by the potentially complex challenges of implementing security controls around non-public information should understand that the attempt is in itself worthwhile.

But let us also recall Schwartz and Janger's second concern that outside authorities are usually not aware of how an organization developed its understanding of what "reasonable" is. During an outside audit, especially one in which an exposure is being analyzed, the term "reasonable" should not be contentious or vague, but well demonstrated and communicated. Organizations that adopt the framework will be using implementation and assessment methods that will help them explicitly demonstrate reasonableness as they define it. It is very important for organizations who implement the framework to also record their decisions in the form of minutes, work papers, sketches and memos for posterity. If for any reason they need to demonstrate or defend the word "reasonable" to any outside authority, evidence from implementation of the framework will go a long way toward that demonstration.

Organizational Commitment

Secure document handling is like insurance. Both are more expensive doing with than doing without – at first. Secure document handling is more costly in staff time and in money. Organizations generally do not want to spend money if they do not have to. So if an organization has decided that it will not commit to secure document handling, then any independent effort by staff or management to secure documents will only be supported by sheer will, time intensive procedures, and will lack enforceability. These may not be enough to mitigate the costs of the inevitable exposure.

To say that exposure is inevitable may seem to be an exaggeration. However, the regular reports of exposures of private information listed at sites, such as the Chronology of Data Breaches at the Privacy Rights Clearinghouse web site⁷ help demonstrate how common exposures are. Keep in mind that sites such as this only list the exposures that are discovered and reported. But as well, secret

government documents roam the Internet freely, according to U.S. Congressional testimony in July, 2007⁸. Documents, such as counter-terrorism strategies, the Pentagon's IP network architecture and IP addresses and many other highly sensitive documents have been found on peer-to-peer networks. If the Pentagon is losing data during war time, it is safe to believe that information exposure is inevitable.

Organization members who need to convince their leadership to adopt document security controls may find it useful to bring to their leaders' attention cases in which other organizations were harmed by losing control of information. Perhaps an internal formal review of document handling that demonstrates how the organization loses control of non-public information is important. If the organization is governed by an audit committee or a compliance officer, these are excellent parties to work with as they are focused on managing risks like those that come with possessing non-public information. Some organizations are required by regulations, laws and industry standards to secure data by specifying what information must be handled and in what way. Gramm-Leach-Bliley, Sarbanes Oxley, HIPAA and the PCI Security Standards Council all provide specific standards and rules for demonstrating the proper secure handling of non-public information. Not surprisingly, auditing the effectiveness of security controls is required by each of these standards and regulations. Appealing to leaders to ensure compliance with such standards, if they apply but have not been implemented, is also critical.

Perhaps the first indication that an organization is committing to data loss prevention is by having an officer of the organization assign responsibility to a team of people to implement and audit compliance with security controls. Such a team must have the time and resources to satisfy their responsibility (anything less would be unreasonable). As well, the team should include people who handle documents, those who administrate documents and operations, and any compliance or internal audit staff if the organization has such functions.

Finally, any appeal to leadership about the need to commit resources and authority to develop and improve document security controls should be met with recommendations. Whether the recommendations come from observed best practices, explicit controls standards or this framework, an organization's leaders will likely struggle with whether they can secure non-public information if they do not see how it can be done.

An excellent feature of the Controlled Event Framework is that it can be, in fact should be, implemented incrementally. Incremental adoption will be described later in the document, but this method of testing controls before applying them universally will appeal to the more cautious of organizational leaders; especially when the attempt is itself more protection against liability than doing nothing at all.

Risk Assessment: Itemizing Risk Categories

When an organization decides that it will invest in document loss prevention, a good business question to ask is how much to invest. In fact, this is our Essential Business Question Number 1: “Are our security controls reasonable?”

To answer the question, the implementation team must determine what its risks are. Risk assessments can range from the thorough to the casual. In the Controlled Event Framework, risk is considered for three different purposes; to create risk categories, (which will be discussed in this section), to create instructions for document handling (using a threat tree analysis), and to measure the effectiveness of controls (which we will explore in the “Auditing the Framework” section). Each of these purposes requires a different level of assessment.

The first of the three risk assessment methods is fairly simple, and is the one we will be discussing in this section. Its purpose is to determine how many risk categories the adopting organization will require in their framework. This can likely be determined by focusing on the impacts that the implementation team has identified.

After creating its lists of threats, vulnerabilities and impacts that effect non-public documents, an implementation team should consider the following questions while developing their risk categories:

1. *Of the impacts we have described, which will likely require a level of security rigor that is distinctly higher or lower than others?* Let’s recall the law firm we described earlier. Most of their business is based on real estate transactions, so most (presumably) of their documents are public. Some of their documents, for corporate mergers, will have lots of market sensitive information, meaning that

documents are sensitive to clients, but not protected by regional laws. The least amount of their work is class-action malpractice. These documents will be protected by regional laws when they contain personal or reputational information. As well, the firm possesses administrative and employment documents.

The firm may initially decide to prioritize its corporate merger information security over its medical malpractice information and internal administrative security, since PII occurs in those documents but in small numbers. And it may prioritize its real estate documents least, since those are public.

2. *In a walk-through of events, do we require distinct handling instructions to protect us from distinct impacts?* While creating risk categories, the implementation team should walk through a series of events to think through the instructions and controls they would require to secure documents during that event.ⁱ It is appropriate to sketch out what those instructions or controls might look like. The firm may realize while walking through the handling events that they will not care about restricting copies of real estate records, but that they want to minimize the number of copies of all other types. This suggests two risk categories: real estate and all other.

But if while walking through the *Dispose* event of their medical malpractice cases they determine that they should never make copies of medical PII, they may consider adding a risk category. They may decide that PII data may not be printed, converted, or copied unless it has been inventoried and controlled. This implies a stricter standard than handling corporate merger documents which must be copied minimally. Now three risk categories become apparent: *real estate*, *corporate mergers* and *medical malpractice*.

The firm's administrative employee records may be considered private enough that they decide to include them in the medical malpractice or corporate merger risk categories.

ⁱ The implementation team has not yet created its definitive list of events. However, they may use the generic list provided in this document as a starting point.

Earlier, we discussed the “California Counties ‘Best Practices’ Information Security Program.”

That document proposes two criteria for classifying information: what is the risk associated with the document, and will information handlers be able to easily determine categories based on some quality of the information. What the law firm should consider important at this point in their risk category development is the naming of the risk categories. If administrative staff members are handling employee documents with PII, what makes them think that they should be handled with the same risk profile as *corporate mergers* or *medical malpractice*? Remember that the primary function of the Controlled Event Framework is to increase awareness among document handlers.

The firm may then decide that all documents, including those used in their real estate work, are managed with a risk category called *Standard Security*, corporate mergers require *Confidential* and employee documents and medical malpractice require *Personal Information*. When an administrator decides at tax time to print payroll records, she can instinctively tell what risk category will give her the necessary handling guidance just by looking at the risk category names.

3. *Now that we have sketched our instructions, can we afford to implement them?* An extremely important part of the framework is determining whether controls are reasonable. While establishing risk categories, an implementation team must determine whether the security rigor implied by a risk category creates handling instructions that can actually be followed. In the case of our example law firm, they will not really know this until they start designing instructions for events in each risk category. If, for instance, they realize that they cannot afford a system that prevents them from making copies of PII - a driver for the *Personal Information* risk category - they may want to change their expectations of how many risk categories they can actually manage to. They may decide to eliminate the *Personal Information* risk category altogether, and include its member documents in the *Confidential* risk category. But again, this determination is likely only considered when instructions for events are being designed.

Finally, the firm will realize that they have created risk categories based on risks as they pertain to their behaviors: at the level of events, of course, but also at the level of the types of work they do. They are asking themselves “What work do we do that exposes us to types of risk?” This is very

different from many document classification methods that ask, “What information is in this document, and who is allowed to view it?”

Operationally, the firm has used the Controlled Event Framework to classify documents according to the type of work they do, which means that any document used in a medical malpractice case - whether or not it contains PII - is handled at the same level of security as PII. The law firm may decide to refine their document handling instructions to make the distinction between individual documents that actually contain or do not contain PII. However, it is also reasonable to assume that in the course of work PII information may show up in any document related to a medical malpractice case. With this concern, the firm may decide that all of the documents related to medical malpractice will be protected by the “Personal Information” risk category. This is a choice that the firm needs to make for themselves. As we saw in the Schwartz and Janger article, the courts are giving organizations broad leeway in determining what is reasonable for them.

Itemizing Events

Earlier in the paper, we reviewed a list of 11 events that organizations will likely need to consider; categorize, receive, store, copy, create, use, backup, send, dispose and archive. Many organizations will have other significant document handling events. The law firm we discussed will need to consider adding Bates stamping and redaction. Engineering firms will need to consider version control. Security analysts, computer forensics professionals and penetration testing organizations may want to consider whether certain tests or exploits on information assets require special events, such as packet capture, recovering files or logging tests.

In order to develop a good list of events, the implementation team will need to have an understanding of everyday work that document handlers and administrators engage in that have some effect on the security of a document. This is an excellent example of why the implementation team should have broad representation from across the organization.

Finally, if the list of events at the beginning of the implementation is not all encompassing, the

team should not be concerned. The framework is extendible and can have events added or removed, depending on the organization's needs.

Itemizing Formats

Remember Essential Business Question Number 2? It asks, "Are our security controls understood?" An important way that the framework handles this question is by making distinctions based on document formats, whether those formats are paper documents, relational database files, spreadsheets, document images, audio recordings, transparencies, digital diagrams or whatever an organization defines as a document containing information it wishes to secure.

When instructions are based on events, the instructions will likely be different for, say, *Copy* instructions for a *Database File* and for *Paper*. If a user is faced with a generic instruction, "The user may only copy documents to a device that they exclusively control while in use," then how does the data analyst know whether they are allowed to export data to a spreadsheet for analysis?

Applying handling instructions at the level of the document format helps an organization know that their employees understand security controls.

This does not preclude an organization from using a generic document format. Having a format *Generic* has four uses. If an organization only cares to control a small variety of documents and plans to handle them in the same way; for example, paper, microfiche and developed film, then a format distinction is not helpful. *Generic* should suffice.

The *Generic* format type also provides excellent benefits when used in conjunction with other formats. It can provide high-level guidance for any instructions to be used during an event. As the implementation team begins to formulate instructions based on events and risk categories, the team can start by creating guiding principles for each event in the *Generic* format type.

For instance, *Receive* instructions for *Generic* can read "Verify that the contained information

matches the risk category. Inventory the documents. Secure delivery media. Transfer documents to a location that matches *Store* requirements.” This provides excellent guidance for developing *Receive* instructions for any format. Boxes of paper that may be received will have different specific inventory instructions than electronic files that arrive via an FTP server, but the implementation team will have common guidance for each event and risk category that the individual file formats will conform to.

By using the *Generic* format type, the implementation team also will have a way to communicate to upper management what the handling standards are for each risk category. While specific handling rules are important for people who handle documents, they are un-necessary and probably overwhelming for those who don't.

Finally, the *Generic* format type provides document handlers with guidance in case they encounter documents that they don't have instructions for. If they normally encounter three format types during their work, then they come into contact with documents of an unexpected file format, then they will know, in principle, how to handle those new documents.

Developing Instructions

The implementation team will create boilerplate instructions for events for each document format and each risk category. Again, this demonstrates the need for the implementation team to have a broad representation of organization members. Experience matters here.

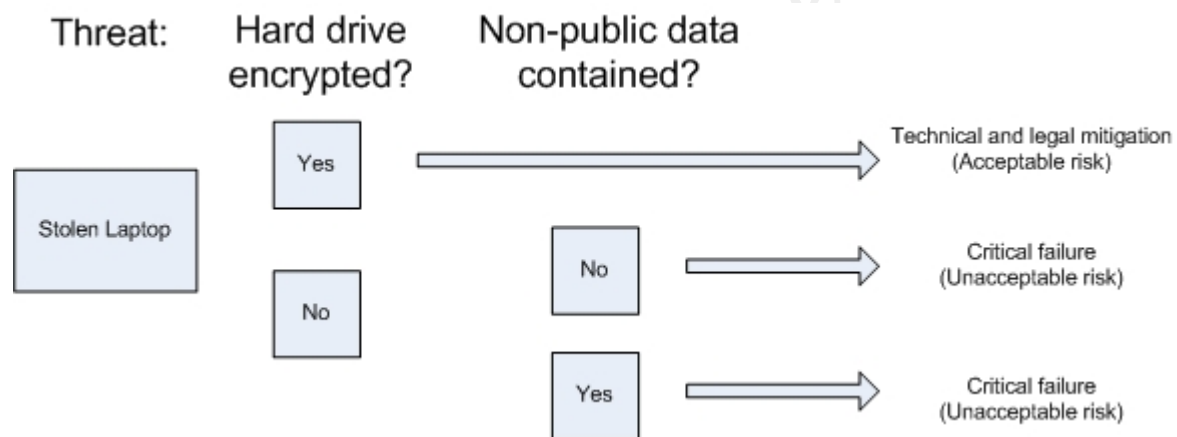
The development of instructions is the method that the Controlled Event Framework uses to ensure that threats to the organization's documents are countered by controls, and that the controls are well known across the organization.

We will demonstrate how instructions can be created using the second of three risk assessment methods as described earlier in this section, using a threat tree (normally known as an “event tree” but renamed here to avoid confusion with the Event framework element).

In Section 3: The Problem, we listed several vulnerabilities and threats that an organization deals with while managing document security. The implementation team must create a list of vulnerabilities and threats that it considers to be reasonable risks, then derive instructions that are designed to reduce the vulnerabilities and to counter the threats.

A worthwhile method for accomplishing this is to create a threat tree, as represented below.

Figure 5



This threat tree chooses a threat to analyze, then lists vulnerabilities horizontally to determine which vulnerabilities need to be controlled. Using this simple analysis, several benefits are gained. The implementation team has:

- Developed an instruction or control, e.g. "Encrypt hard drives."
- Demonstrated that they considered how to mitigate known vulnerabilities when they created their handling instructions
- Determined what is and what is not an acceptable risk when vulnerabilities are matched with threats
- Created an artifact for auditors to determine the impact of controls that tested as failures
- Created an artifact of its reasoning for determining what was reasonable as a security

control.

The implementation team will want to run this exercise for as many threats as they find likely. A good source of threats to develop their list can be found in an unsettling but important web site, the Chronology of Data Breaches. The site which is maintained by the Privacy Rights Clearinghouse (referenced in the end-notes) lists all publicly announced breaches of personal information in the United States recorded since January, 2005. The list describes who possessed the data that was exposed and how much data was involved. But more importantly for this effort, it described how the data was exposed. A perusal of this information will provide a useful list of likely threats for the implementation team's threat list.

But the implementation team must also look closely at their own organization for vulnerabilities to threats. A thorough third-party penetration test would be beneficial here, especially if the penetration testing team focuses on systems and events that handle non-public documents.

The implementation team should not be too concerned about ensuring that the instructions are perfect when they are first written. In fact, the adopting organization should implement the instructions incrementally and develop them after trial and error. This will also be key to demonstrating whether controls are "reasonable." We will discuss incremental adoption at the end of this section.

Handling instructions, as we have stated earlier, should take the form of "Who can do what and where?" An example we used earlier was related to an engineering environment: "The lead engineer or the lead engineer's designee must shred all super-ceded draft designs using the e-doc file shredder and store all shred reports in the design plan log."

The implementation team should decide the level of detail they want to include in the instructions that are contained in the framework. Detailed rules that the engineers must follow while shredding files may best be described in policies and procedures, while leaving the essential instructions in the framework. The above instructions may better read as follows: "Approved engineers must shred all super-ceded draft designs and store all shred reports as described in the 'Shred Electronic

Files' procedure." These instructions provide the document handlers with enough guidance to distinguish between shredding a file and simply "deleting" it, while also telling them where to get more detail in case they need to be refreshed on specific requirements. In this way, the document handlers get the guidance they need, but not so much that they are overwhelmed or annoyed.

Until now, we have focused on how instructions are useful for people who handle documents. However, they are also essential for administrators and auditors. When we have established instructions for handling documents, do office support and IT staff know what is required? Are administrators providing the support that document handlers need in order to satisfy the handling requirements?

So while we have talked about instructions in terms of what users may, must or may not do, they also must tell administrators what is required of them, and must tell auditors how to determine if the instructions were followed. For this purpose, three sets of instructions should be considered for each combination of events, risk categories and formats.

The instructions for administrators should be developed with the administrators. Their instructions will rely on controls, whether technical or procedural, that may or may not exist or that may or may not be easy to change.

The instructions for auditors should take the form of test steps that would determine whether the user and administrator instructions were followed.

If this sounds demanding and overly complex, it will all become very easy to envision and manage in the following section, "Database Application."

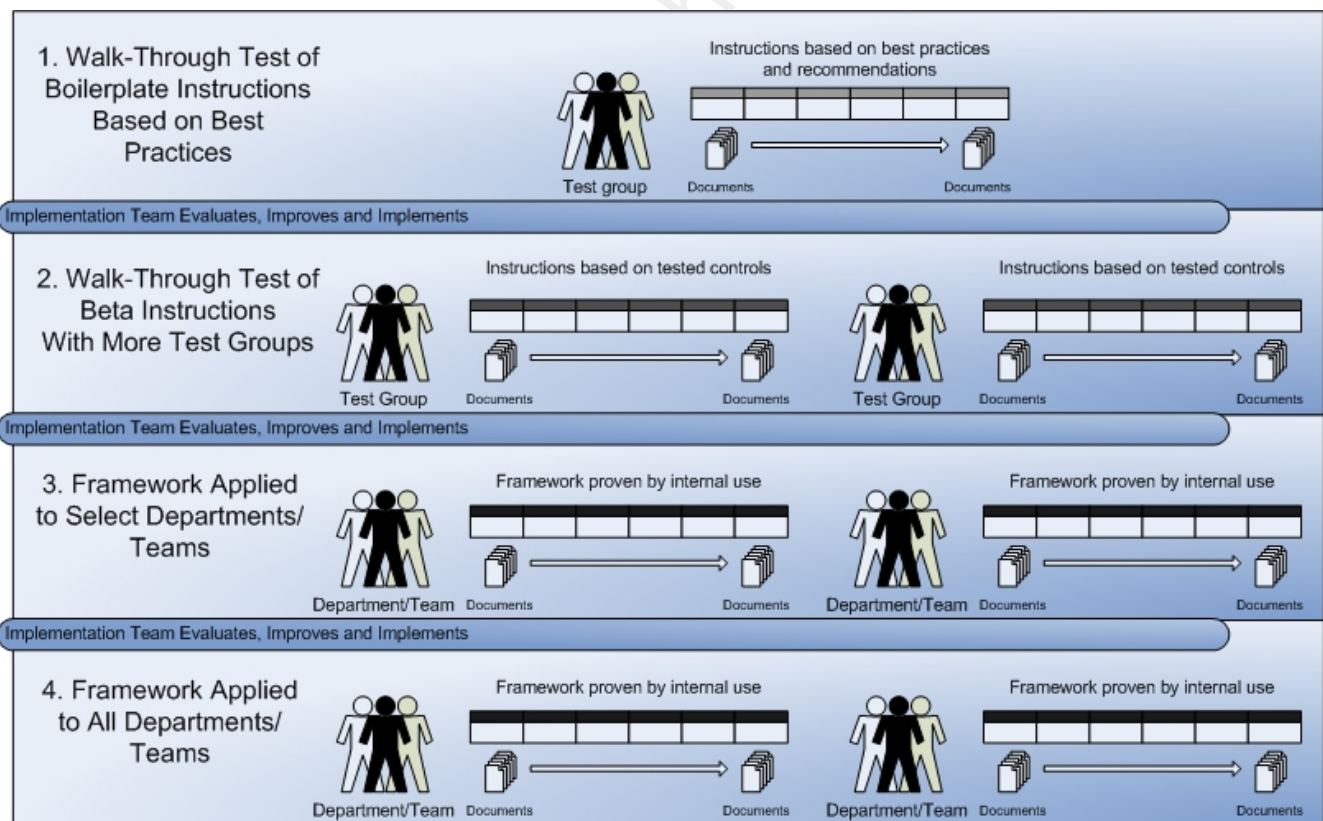
Incremental Adoption

An excellent feature of the Controlled Event Framework is that it can be, in fact should be, implemented incrementally. After draft instructions are created, they can be tested to a limited set of documents, or in a limited population of users. Where instructions are deemed to be beneficial and

effective, they can be adopted as the final rules. Where the instructions are deemed too weak after testing, they should be strengthened. Where they test as unsustainable or too expensive to function properly, they must be scaled back. Incremental adoption is yet another process with which that the organization demonstrates why its controls are reasonable.

Figure 6 represents the incremental adoption approach, displaying how an organization uses the approach to test the reasonableness of instructions. And while the approach is important during implementation, an essential aspect of the framework is to continually test for compliance during its use. Improvements and adjustments to controls and instructions will continue during the lifetime of the framework.

Figure 6



6. Database Application

Why an Application?

When we talk about a framework, we describe a model for representing and managing complex systems. Perhaps the most common way for organizations to implement frameworks is through database applications. In fact, the many components to the Controlled Event Framework; its elements, implementation and testing, are so complex when combined that it is difficult to see how such a model for document loss prevention could be effective without a database application to support it.

In this section, we will discuss the basic elements of a database application that could support the Controlled Event Framework. The intentions of this section are to demonstrate a valuable method for implementing the framework, and to help the adopting organization to actually see how all of the components of the framework fit together from concept to delivery to testing.

The Tables

The most basic components of the database structure will be the tables. The first of these tables is the Risk Categories table seen below:

Figure 7

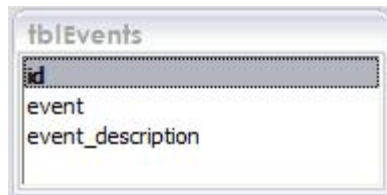


tblRiskCategories	
id	
category_name	
category_description	

The table is simple, allowing the adopting organization to list and describe the risk categories they will be using in the framework.

The next table is for Events:

Figure 8

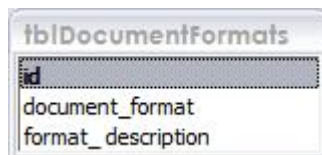


tblEvents	
id	
event	
event_description	

Again, we have a simple table that lists the events and their descriptions.

The table for document formats as just as basic:

Figure 9



tblDocumentFormats	
id	
document_format	
format_description	

In this table, we have a list of document formats, and their descriptions.

Next, we have the Event Framework table which brings the previous three elements together, and adds the handling instructions, and essential testing requirements.

Figure 10



tblEventFramework	
id	
risk_classification	
document_format	
event	
user_instructions	
test_instructions	
admin_instructions	
implemented_status	
incident_response_units	
incident_detect_time	
incident_response_time	
incident_protection_time	

The first three fields after the ID identify a record by uniting a risk classification with a document format and an event.

By looking at these three fields in the populated table, you can expect to see a basic record set

that looks like this:

Figure 11

<i>risk_category</i>	<i>document_format</i>	<i>event</i>
Standard Security	Paper	Receive
Standard Security	Paper	Store
Standard Security	Paper	Copy
Standard Security	Paper	Dispose
Standard Security	Editable e-Doc	Receive
Standard Security	Editable e-Doc	Store
Standard Security	Editable e-Doc	Copy
Standard Security	Editable e-Doc	Dispose
Personal Information	Paper	Receive
Personal Information	Paper	Store
Personal Information	Paper	Copy
Personal Information	Paper	Dispose
Personal Information	Editable e-Doc	Receive
Personal Information	Editable e-Doc	Store
Personal Information	Editable e-Doc	Copy
Personal Information	Editable e-Doc	Dispose

Each record will also contain the handling instructions for users (*user_instructions*), administrators (*admin_instructions*) and auditors (*test_instructions*).

Also notice the field that allows the implementation team to state whether the instructions they list are implemented or not (*implemented_status*).

The implementation team will populate this table with every possible combination of Risk Category, Document Format and Event. If the implementation team has listed four risk categories, six document formats and ten events, they will have 240 records in this table ($4 \times 6 \times 10 = 240$). Each record will also store the instructions for users, administrators and auditors in the appropriate fields.

The instructions in each record should be considered as the minimum security instructions allowed for that risk category and document format combination. While describing the Issued Instructions table below, we will see the role this table plays while the organization issues the actual instructions to its document handlers.

Time Based Security at the Event

This table also contains four fields that are associated with incident security metrics; Protection Time, Detection Time and Response Time. In the section “Implementing the Framework” we discussed risk assessment and mentioned that there are three purposes for risk assessment. The first is used to establish the risk management categories that would be used in the framework. The second was used to create instructions that countered known risks. The third risk assessment method is meant to measure the effectiveness of controls.

The fields *incident_detect_time*, *incident_reponse_time* and *incident_protection_time* contain numerical values that quantify how long it takes the organization to recognize a violation of the instructions and, perhaps, an exposure of the information. The field *incident_response_units* simply tells us whether the numerical values represent minutes, seconds, days or months.

To demonstrate why these fields are useful, let’s first understand the idea of time based security. The basic idea can be expressed simply as “Protection Time must be greater than the Detect Time plus the Response Time,” or $P > D + R$

In this equation, “Protection” references controls that provide security for a document for longer than it would take for people to “Detect” a violation and “Respond” to it.

The purpose of the Controlled Event Framework’s security instructions is to protect an organization’s documents. It does so in part by using these time-based security metrics. If a database file was copied to a local machine despite instructions to not do so and a monthly audit would detect the disallowed copy, then the detect time of the violation would be, worst case, 31 days. If the time required to remove the unauthorized copy is one day, then the total exposure is 32 days. Protection, then, should provide security of database files for at least 33 days.

In the framework, the description of the Protection parameter lies within the *admin_instructions* field. In this case, the instructions could read “Only DBAs are allowed to copy database files from their production locations, and only by following quarterly maintenance procedures. Database files are stored in directories that allow only database administrators and necessary services to interact with the files at the file level. Event logs must detect and alert when permissions on the directory change, and when file access rules are violated.”

The administrator’s instructions now tell us Protection equals about 91 days. How do we know this? The access rules prevent anyone from accessing the database files, but for the DBA. The DBA may make copies of the database files every quarter (91 days) during maintenance. Because the administrator can place a copy of the file in an unauthorized location, then Protection can only be said to work for 91 days. Given that the exposure (Detect + Respond) was 31 days + 1 day, we would see that the time based security equation indicates well-designed security controls: $91 > 31 + 1$.

The benefit of time based security in the Controlled Event Framework is that it provides risk assessment at the level of the event. It helps an organization answer Essential Business Questions Number 1 and Number 3, “Are our security controls reasonable?” and “Are our security controls working?”

Let’s imagine that the organization was uncomfortable with the risk that a DBA could, on a quarterly basis, make an unauthorized copy of a database file. Let’s say that they are considering investing in a system that would detect and alert when database files are being copied to an unauthorized location. Now let’s imagine that such a system costs tens of thousands of dollars, and would take six months to implement. This new system would provide, they calculate, a protection time of 365 days. Is the security control reasonable? Are 91 days of protection not sufficient? The organization can answer these questions based on the risk they have established, and the known alternatives.

If the organization can calculate the cost of exposing a database file for 274 days (365 days of potential protection – 91 days of current protection) then it can also, using time based security, calculate whether the cost of 274 additional days of protection is worth their expense. This will take a mature risk assessment process to calculate. However, time based security provides a consistent method for demonstrating whether controls that are in place and controls that were considered, are reasonable.

In the section “Auditing the Framework” we will describe how time based security is considered during the testing of the framework.

Issued Instructions: Flexibility in the Framework

One last table we should discuss is the Issued Instructions table. Issued Instructions are a record of every occasion in which an instruction set was created within an organization.

Figure 12



id
matter_ID
risk_category
document_format
description
unit_description
count
variability
classify
receive
store
use
convert
copy
print
send
backup
archive
destroy

The Issued Instructions table represented above would be used by the example law firm we

have described already. They issue instructions for each “matter,” or engagement, they work on. (An engineering firm may not use “matter” but rather “project” and “version” fields. A retailer may substitute the matter_id field with “department” and “division” fields.)

Each row in this table reads like the instructions for handling a document format that is being used in the matter. If one matter uses two document formats, the matter will have two rows in this table. Each row in the table reads, “This engagement will use ‘X’ document format within ‘Y’ risk category.” This pairing of document format and risk category can be made in an application form that relates the matter_ID with the Event Framework table. Then using a stored procedure, the record’s event fields (classify, receive, store, use, convert, copy, print, send, backup, archive, destroy) are populated with the instructions found in the Event Framework table for the pairing of the formats and risk category.

The experienced database engineer will ask why the events fields are used in the Issued Instructions table if events are already stored in the Event Framework table. Perhaps the best way to think of the Event Framework table is as a look-up table. While issuing instructions the Event Framework table provides minimum handling requirements that can be over-ridden in the Issued Instructions table given special requirements at issuance. If the organization needs to *ad hoc* increase the security rigor for a specific event while instructions are being issued, they should have the flexibility to do so, but must record their custom instructions.

For instance, in our example law firm, a partner in the real estate practice knows that the documents she is using in her matters are public, and are being handled with a set of effective security instructions. Even though the standard instructions allow it, she considers it too risky to have backup copies or archives of documents used in her matters. So she requests, when a new matter begins, to issue instructions to her associates and the firm’s administrators that all Standard Security instructions apply, but to make no backups and to make no archives of her engagements’ files.

If the Issued Instructions table relied on the standard handling instructions contained within the Event Framework table, the vigilant partner would have a more difficult time communicating her requirements to all involved parties, including the auditor whose job it is to verify that her security instructions were followed.

While this feature may allow *ad hoc* decreased scrutiny on an event while issuing instructions, the database application should have a rule or approval procedure preventing such customizations.

Issued Instructions: The Instruction Set

Despite the complexity and detail that the Controlled Event Framework seems to impose on the adopting organization, it allows organizations to create very clear and targeted instructions to the individuals who use documents, administer documents or audit organizations for their secure handling of documents.

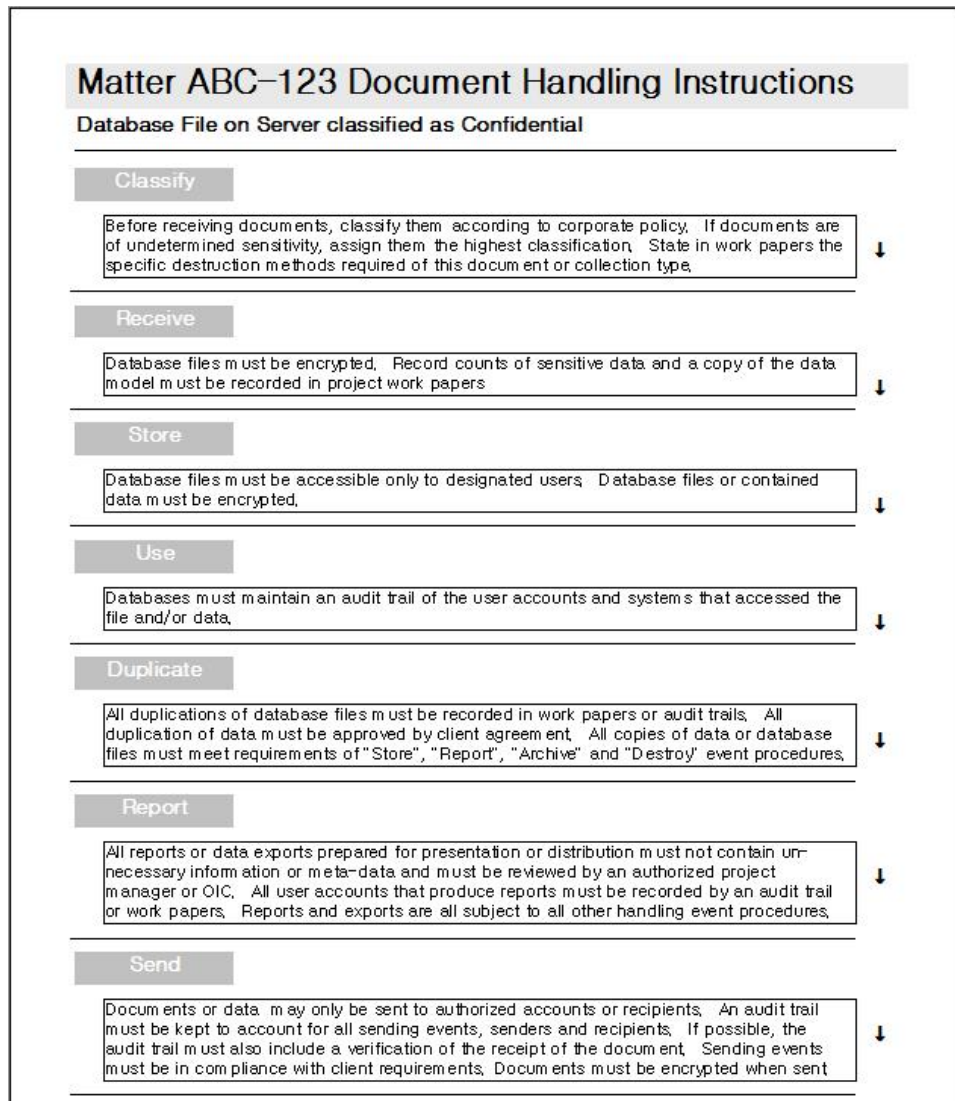
When a new matter begins, an administrator at the example law firm may use an application form in the database to create records in the Issued Instructions table. The matter may have the following requirements: We are starting a case with confidential information. Not personal data, but confidential. The documents will be: a database file that will sit on our server; twelve boxes of documents, already Bates stamped; and electronic images of e-mail. The database file may only be copied with the client's permission.

The application form then looks inside the Risk Categories table, the Document Formats table, the Events table and the Event Framework table. It selects the appropriate records from those tables and populates the Issued Instructions table with three new records, one for each format.

The administrator can then adjust the new instructions record for the database file to require that all copies be approved by the client.

At that point, the administrator can print the instructions that the partner will distribute to her associates. The instructions will look something like this:

Figure 13



Instructions for administrators and auditors can be similarly produced.

7. Auditing with the Framework

Organizations that handle non-public information are often required to handle business processes, data and documents under scrutiny from outside authorities. Law makers, certification

Chris Cronin

45

bodies and regulatory agencies issue requirements to organizations for secure handling of documents and data, whether they are state laws, the Securities and Exchange Commission regulations, or professional standards bodies, such as the PCAOB and the PCI Security Standards Council. When auditors test compliance with these outside requirements, they can do so by using the guidelines that these requirements provide. In fact a key benefit to using a framework for implementing controls is that outside authorities can quickly understand how an organization manages its risk.

And while many of these organizations overlap in controls requirements, there are many business processes and organizations that may not be subject to these rules. For instance, Sarbanes-Oxley rules do not require that sensitive documents are protected unless they are related to financial reporting.

The Controlled Event Framework for Information Asset Security provides a method for organizations to determine and manage their risk in handling non-public information, but it as well provides a necessary mechanism for communicating with outside authorities and auditors.

Schwartz and Janger in their article found that organizations lack a mechanism for communicating how they determined their risk. If they cannot communicate that with auditors, whether they are from a certifying body or a plaintiff, then their risk is heightened. Working within a framework to demonstrate how a risk assessment drove the development of reasonable controls is critical to successfully communicating with auditors.

Let us also recall Essential Business Question Number 3 which asks, “Are our security controls working?” An organization cannot know that unless they audit their controls. Moreover, if an organization does not know whether their document security controls are working, then they become subject to the Arthur Andersen rule: Document handling policies are an added risk. If that risk is not managed by verification or auditing, then it results in only occasional adherence to policies. This can be seen as capricious. In the wrong instance, an organization can pay an unbearable price for that appearance.

In the case of the Controlled Event Framework, the framework and auditors support each other.

While the auditor reduces the adopting organization's risk of becoming victim to the Arthur Andersen rule, the framework provides practical information for the auditor to conduct their testing.

As we described the creation of instructions in the "Implementing the Framework" section of the paper, we mentioned three activities that feed the auditor's work; providing test instructions to auditors, providing time-based security metrics, and using a threat tree to analyze risk. So while an organization's use of the Controlled Event Framework makes a useful audit possible, it is these three activities that provide the most practical benefit to the auditor.

The auditor's tests should verify that user and administrator instructions were being followed, that the administrator's tools (usually systems controls and policies and procedures) were validated through penetration testing, and that time-based security metrics test effectively.

As well, audit tests and results should be recorded in detail. A record of regular internal audits and their findings is essential protection against the Arthur Andersen rule. An internal auditor can manage their audit work within a table that is used to store audit findings. Such a table should record the findings of an audit associated with a scope that is meaningful to the organization (in the case of the law firm, the test can focus on a specific matter), with tests that are associated with events.

The table below represents what the law firm's internal audit department may use to demonstrate the effectiveness of controls, and to record that the controls were being assessed in case this oversight needs to be demonstrated to outside authorities.

Figure 14

Matter_ID	Format	Event	Instructions	Date Tested	Observations	TBS Test
-----------	--------	-------	--------------	-------------	--------------	----------

ABC-123	Database file	Classify	Before receiving documents, classify them according to corporate policy. If documents are of undetermined sensitivity, assign them the highest classification. State in work papers the specific destruction methods required of this document or collection.	6/1/2001	Passed. Documents were classified as required. Information appears to qualify as "Confidential"	None required
ABC-123	Database file	Receive	Database files must be encrypted. Record counts of sensitive data and a copy of the data model must be recorded in project work papers.	6/1/2001	Passed. Files were encrypted as received. Inventory documents match record counts. Encryption method tested as effective.	None required
ABC-123	Database file	Store	Database files must be accessible only to designated users. Database files or contained data must be encrypted.	6/1/2001	Passed. File folder had permissions only for designated attorneys and DBAs. Audit logs showed appropriate access. Penetration test from previous quarter verified	Successful

					effective access controls at the file server.	
ABC-123	Database file	Use	Databases must maintain an audit trail of the user accounts and systems that accessed the file and/or data.	6/1/2001	Failed. DBA audit logs not functioning. No record of user access was recorded	Failed. System alerts that were designed to warn on audit logs did not function.
ABC-123	Database file	Duplicate	All duplications of database files must be recorded in work papers or audit trails. All duplication of data must be approved by client agreement. All copies of data or database files must meet requirements of "Store", "Report", "Archive" and "Destroy" events.	6/1/2001	Failed. DBA audit logs not functioning. No record of user queries was recorded	Failed. System alerts that were designed to warn on audit logs did not function.
ABC-123	Database file	Report	All reports or data exports prepared for presentation or distribution must not contain un-necessary information or meta-data and must be	6/1/2001	Failed. DBA audit logs not functioning. No record of user queries was recorded	Failed. System alerts that were designed to warn on audit logs did not function.

			reviewed by an authorized project manager or OIC. All user accounts that produce reports must be recorded by an audit trail.			
ABC-123	Database file	Send	Documents or data may only be sent to authorized accounts or recipients. An audit trail must be kept to account for all sending events, senders and recipients. If possible, the audit trail must also include a verification of the receipt of the documents.	6/1/2001	Failed. E-mail DLP services not yet deployed. No evidence gathered. Security department projects 45 day roll-out.	Failed. E-mail DLP services not yet deployed.
ABC-123	Database file	Archive	Client requested no archiving.	6/1/2001	Passed. No record on file with third party archive service to archive records for this matter. Archive service has current SAS70 Type II so their record keeping is reliable.	Not tested

ABC-123	Database file	Dispose	The destruction of documents must be recorded to include the account or user that destroyed the document or file, the time and date of the destruction, and the method of destruction.	6/1/2001	Matter still active. No dispose requirements were triggered.	
---------	---------------	---------	--	----------	--	--

This audit table records essential information for the audit, namely:

- The audit test scope was for each engagement and format
- The provided instructions (we show just user instructions due to space constraints, but administrator and test instructions will also be relevant)
- When the test occurred
- The observation of compliance with the instructions (which declares a pass or fail and evidence for the observation)
- And the observation of compliance with a time based security metric (TBS), if it is relevant.

Audit tests should rely on the instructions that are stored within the framework. However, it must also ensure that penetration testing has been conducted and that time based security metrics were verified through walk-through tests.

Time based security and penetration tests provide auditors with a measurable way of determining if administrative instructions and controls are well designed, but a test of the instructions should also determine two other issues, whether the user and administrator instructions were followed, and whether the administrative controls protect information assets as designed.

Looking back at our audit table example in Figure 14, we see findings that result from two system failures: the audit logs failed to work on the database engine that hosted the database file, and

an automated Document Loss Prevention application that prevents sensitive e-mails and file transfers from being sent is not yet operating. A series of failures occur in the audit, but only two causes are identified.

Now imagine that over several months, the DLP system is never installed correctly. The security team may be able to demonstrate that they are too committed to other controls and have not had time to implement the DLP. Or their DLP product just did not work as advertised. A history of detailed audit reports will demonstrate the history of failures of the controls, and indicate that the control as originally designed was not reasonable for that organization.

Audits are useful not only to determine whether controls work and instructions are being followed, they also are critical to the organization for understanding if they have created expectations that they could not meet. But most critically, they prevent an organization from being subject to the fate of Arthur Andersen.

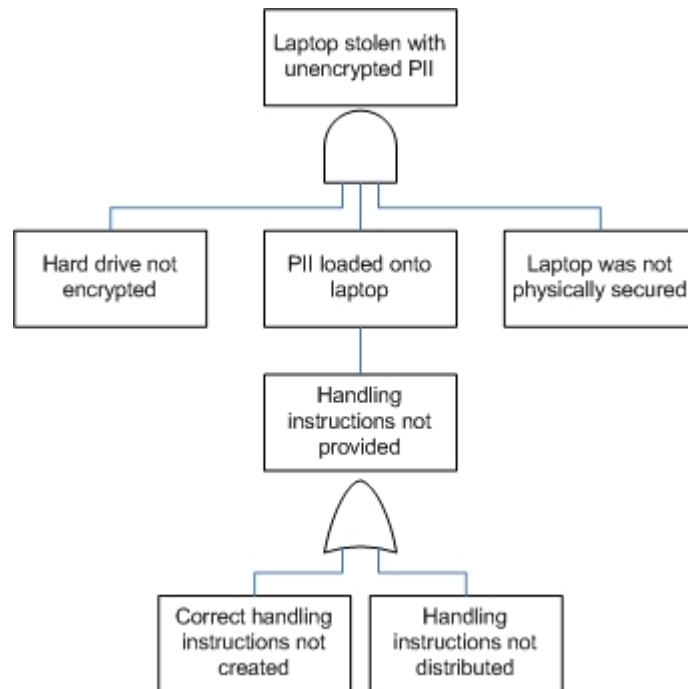
Finally, the auditor must report their findings to the managers who are responsible for following instructions, whether those are administrative managers, or the managers who oversee the document handlers.

The findings must be met with recommendations for improving controls. If an organization finds that its controls are not working, then are they reasonable? Possibly not. So how does an auditor make recommendations that are reasonable? Recall that the implementation team developed boilerplate instructions by using a threat tree to determine how threats should be countered. An auditor can come to her or his recommendations by using an analysis similar to the threat tree.

In the case of auditing, when an adverse event occurs, or a control failure is noted, the auditor has an opportunity to determine the cause of the event. In this case, a fault tree analysis can be used.

A fault tree is a diagram that determines what caused an adverse event or control failure. We will demonstrate a fault tree using the same threat that we used in the threat tree example: a stolen laptop.

Figure 15



The elements of a fault tree that are in use here are the adverse events and control failures (the rectangles), the “AND gate” (the arch that links the adverse event to co-existing failures) and the “OR gate” (the rounded triangle, representing possible control failures that could have led to the preceding failure).

When an auditor provides recommendations in the audit report’s findings, the recommendations should be based on an analysis of root causes. This is a useful way for an auditor to document their analysis, and to evidence the role an audit plays in determining why controls or instructions are reasonable.

The Auditing Limits of the Controlled Event Framework

The Controlled Event Framework creates instructions for users, administrators and auditors

about how documents should be handled at significant moments in the documents' lifecycle. In this way, it is meant to provide organizations with a mechanism for addressing security awareness, designing its security requirements and implementing reasonable controls. Because framework auditing is driven from the framework itself, it relies on other audit procedures to demonstrate the full security of non-public information: namely penetration testing.

Data loss prevention relies significantly on IT systems and procedures, physical property management, and document handling controls. The Controlled event framework addresses the third of these. Therefore, companion audits must be conducted to address the first two. Without a penetration test that attempts to evade the social, manual, physical and automated controls that the implemented framework relies on, critical aspects of the organization's vulnerabilities will be missed, and the value of the instructions provided by the framework will be undermined.

8. Conclusion

Organizations are challenged with myriad regulations, laws, certification challenges and business risks that make handling of non-public information seem daunting. Not only are requirements complex, but the reality of managing the wide variety of document formats makes data loss prevention seem to be an insurmountable process. Even before implementing commercial DLP solutions, these complex requirements must be defined.

With the rise in litigation and increased government requirements over information privacy and security, organizations that take on data loss prevention efforts must contend both with controlling their documents and data and preparing to explain themselves to outside authorities.

The Controlled Event Framework for Information Asset Security provides a way to define the risks that an organization faces, and to provide its members with the instructions they need to reasonably reduce those risks. As well, the framework provides a mechanism for oversight to ensure that the rules are effective, and easy to communicate to management, customers and constituents, or to

outside authorities.

Early in the document, we asked the Three Essential Business Questions:

1. Are our security controls reasonable?
2. Are our security controls understood?
3. Are our security controls working?

The Controlled Event Framework for Information Asset Security answers these questions in several ways.

Are our security controls reasonable?

While “reasonable” is a legal term with much leeway, an organization that handles non-public information must have some way of demonstrating to itself, its customers and constituents, or outside authorities that its controls are reasonable.

The Controlled Event Framework provides answers to this question through:

1. Risk assessments based on impacts to create risk categories
2. Risk assessments based on threats and vulnerabilities to create instructions for handling documents
3. Risk assessments based on measurable effectiveness of controls using time based security
4. Records from the implementation team’s discussions to show how they determined whether instructions were reasonable
5. Records and reports from audits to indicate trends in case an organization must

determine that its attempts at controls were futile, and therefore unreasonable

Are our security controls understood?

A critical component of the Controlled Event Framework is to increase awareness among people who handle documents. Security professionals agree that awareness is key to effective security.

The Controlled Event Framework provides awareness by:

1. Delivering to individuals the instructions for using the documents they handle to the level of security rigor that is required by their content
2. Delivering to administrators the instructions and metrics for securing documents
3. Helping an organization understand its risks by driving three types of risk assessment

Are our controls working?

Finally, in order to avoid the Arthur Andersen rule, an organization needs to know if staff members are actually following the instructions they are provided.

The Controlled Event Framework demonstrates how well controls are working by:

1. Providing an audit function that measures compliance and makes recommendations for improvements
2. Provides a method for determining whether controls and instructions that are never followed are reasonable
3. Delivering to auditors the instructions that users and administrators were presented, and guidance for testing compliance with those instructions.

Organizations face risks not only in handing non-public information but also in communicating

their controls and compliance. A framework as comprehensive as the Controlled Event Framework for Information Asset Security is required to mitigate those risks.

9. Glossary

Term	Definition	Example
Arthur Andersen Rule	Document handling policies are an added risk to organizations.	If document handling rules are not regularly overseen and evenly enforced, occasional compliance may appear to be capricious.
Classification	The required event in which a document or document set is placed into a risk category, based on the risk associated with the information contained within it.	Classification instructions may require the person creating or sending the information to declare the sensitivity of the information before it is introduced into the organization. Classification should be verified at <i>Receive</i> events.
Control	A rule, policy, procedure of a system that reduces the likelihood of a threat.	<i>Encrypted laptop hard drives, security audits, aware staff</i> are controls that reduce the likelihood of a threat.
Data	A series of regularly formatted information stored as records.	<i>Database record sets, spreadsheets, well-formatted delineated files, xml files, packet captures</i> , etc.
Document	An object, electronic file or other set of text, images or recordings that contains or constitutes information.	This can apply to <i>printed paper, photographs, database files, electronic documents, audio or video recordings, packets of data</i> , etc.
Events	An element of the Controlled Event Framework, these are predictable moments in which documents are handled.	<i>Classify, Receive, Store, Copy, Convert, Print, Dispose, Create, Archive, Backup, Send</i> .
Fault Tree	A diagram that displays the cause of adverse events.	
Framework	A normative model used as guidance for managing complex systems.	CobiT, COSO, ISO-27000 series, FISMA.

Instructions: Administrator	Instructions are an element of the Controlled Event Framework. Administrator instructions are explicit rules provided to administrators (staff support or IT) that describe the means for supporting user instructions.	Must contain instructions for handling documents that can be enforced by managing controls. For instance, if files may not be copied, IT administrators must have a system that can enforce such rules. If paper must be stored in locked locations, office managers must provide the facilities for this control.
Instructions: Test	Instructions are an element of the Controlled Event Framework. Test instructions guide auditors so that their testing is based on instructions and controls.	Must provide a test for each instruction and control, including controls that administrators rely on. Some tests may simply refer to penetration tests that validate the security controls that are relied on.
Instructions: User	Instructions are an element of the Controlled Event Framework. User instructions are explicit rules provided to document users that say who may do what and where to the documents they are using.	Must contain instructions of <i>who</i> may do <i>what</i> and <i>where</i> . Required for every combination of an Event, Format and Risk Category.
Risk	A vulnerability when matched with a threat and an impact.	While there are various ways to calculate risk as a value, organizations that consider security should calculate the cost of an adverse event so they know what to invest in order to mitigate the event.
Risk Category	An element of the Controlled Event Framework, this classification of information imposes a security rigor on the instructions provided to document handlers, administrators and auditors. Risk categories should be aligned with risks that the organization has identified.	<i>Low, Medium and High</i> likely have little meaning. An accounting firm will likely consider <i>Administrative</i> and <i>Client</i> as their two categories since all client documents will have the same high legal and business risk and will always outweigh internal administrative documents in the need for security rigor.

Threat	An event that causes an organization to lose control of an asset.	<i>Theft, Accidental loss, Posting on public servers, Non-compliance, Information in unsecured devices</i> are all likely threats.
Threat Tree (Event Tree)	A diagram that displays the impact of threats to vulnerabilities.	
Time Based Security	A way to calculate security controls based on time. If the time that protection is effective is longer than the time it takes to detect and respond to a violation, then protection is adequate.	$P > D + R$

10. References

¹ Greenberg, Pam & National Conference of State Legislatures (January 25, 2008). *State Security Breach Notification Laws*. Retrieved January 31, 2008, from <http://www.ncsl.org/programs/lis/cip/priv/breachlaws.htm>.

² Ponemon Institute, LLC. (2006) U.S. Survey: Confidential Data at Risk. Elk Rapids, MI

³ Arthur Andersen LLP v. United States 544 US 696 (2005)

⁴ Roiter, Neil (May 31, 2007) Springing Leaks: Getting Smart About Data Loss Prevention. *Information Security Magazine*. Retrieved February 2, 2008, from http://searchsecurity.techtarget.com/news/article/0,289142,sid14_gci1256804,00.html

⁵ California County Information Services Directors Association. (2002). *California Counties "Best Practices" Information Security Program*. Retrieved January 16, 2008, from, <http://www.csoonline.com/counsel/kdickeydocument.doc>

⁶ Schwartz, Paul M. and Janger, Edward J., "Notification of Data Security Breaches". *Michigan Law Review*, Vol. 105, p. 913, 2007

⁷ Chronology of Data Breaches. <http://www.privacyrights.org/ar/ChronDataBreaches.htm>

⁸ Boback, Robert, (July 24, 1007). Testimony Before the House Committee on Oversight and Government Reform