Security Risk Analysis and Management: an Overview

Save to myBoK

This practice brief has been updated. See the latest version here. This version is made available for historical purposes only.

The HIPAA security rule requires every covered entity (CE) to conduct a risk analysis to determine security risks and implement measures "to sufficiently reduce those risks and vulnerabilities to a reasonable and appropriate level."[1] The concept of risk management is not new to healthcare. But making decisions about how to comply with a regulation using a risk-based approach is new.

In addition, deciding and documenting the appropriate level of controls based on potential threats that might exploit vulnerabilities is not a customary practice. But the preamble to the security rule makes it clear that regulators recognize that "use of electronic technology...results in many new and potentially large risks. These risks represent expected costs, both monetary and social. Leaving risk assessment up to individual entities will minimize the impact and ensure that security effort is proportional to security risk."[2]

This practice brief reviews the regulatory requirements for security risk analysis and management, provides an overview of the types of risk analysis that can be performed, and offers a practical approach on how to comply with these requirements.

Regulatory Requirement

The security rule requires a CE, in accordance with the security standards general rules (§164.306), to have a security management process in place "to implement policies and procedures to prevent, detect, contain, and correct security violations."[3]

The security standards general rules include general requirements to:

- Ensure the confidentiality, integrity, and availability of all electronic protected health information the CE creates, receives, maintains, or transmits
- Protect against any reasonably anticipated threats or hazards to the security or integrity of such information
- Protect against any reasonably anticipated uses or disclosures of such information that are not permitted or required under the privacy rule
- Ensure compliance with this subpart by its work force

The standards are flexible in regards to approach:

- CEs may use any security measures that allow the CE to reasonably and appropriately implement the standards and implementation specifications as specified in this subpart
- In deciding which security measures to use, a CE must take into account the following factors:
  - The size, complexity, and capabilities of the CE
  - The CE's technical infrastructure, hardware, and software security capabilities
  - The costs of security measures
  - The probability and criticality of potential risks to electronic protected health information

In applying flexibility, however, the preamble to the security rule states, "Cost is not meant to free covered entities from this [adequate security measures] responsibility."[4]

Risk Analysis Approaches

Risk analysis and risk management are two of the required implementation specifications within the security management process standard. The security rule does not specify exactly how a risk analysis should be conducted, but it does reference the

National Institute of Standards and Technology (NIST) Special Publication 800-30, "Risk Management Guide for Information Technology Systems."[5] The NIST publication offers a comprehensive approach to incorporating risk management into the system development life cycle. Threats in the environment are identified, then vulnerabilities in information systems are assessed. Threats are then matched to vulnerabilities to describe risk.

The NIST document includes a description of the roles of various persons in risk analysis and management. It emphasizes the key role senior management plays in understanding security risk, establishing direction, and supplying resources. HIPAA requires assigning responsibility to the security official for the development and implementation of security policies and procedures. This individual may lead the team that actually performs the risk analysis, do much of the policy and procedure writing, and recommend or even select many of the controls.

The fact that NIST identifies the chief information officer, system and information owners, business and functional managers, information technology (IT) security analysts, and trainers recognizes the importance of a team that extends beyond IT and encompasses users.[6] In a clinical setting, users not only can assist in providing application and data criticality information, but must also be involved in determining which mitigation strategies will work.

Because healthcare has so many regulations to comply with, it is also helpful to "piggyback" security onto other functions where possible. For example, a safety officer may make rounds that could include the review of workstation locations. Physical plant security personnel (also known as protective services personnel) may already be monitoring that members of the work force are wearing badges and checking that doors are locked. Communicating these issues with the information security official is critical to compliance.

Qualitative Approach

The NIST approach is generally considered qualitative, because it relies heavily on narrative descriptions of risk. The NIST approach addresses cost/benefit analysis, but not as an integral determinant of risk. Other security experts offer methodologies that rank risks or are highly quantitative.

Ranking Approach

Because the security rule's flexible approach calls for the identification of probability and criticality of potential risks as well as cost considerations, a ranking approach may provide compliance justification and be more conducive to budgeting than a qualitative approach alone.

In a ranking approach, each vulnerability/threat pair can be rated as high-medium-low on a probability scale and a criticality scale. Together the ratings combine scores that can be used to prioritize the risks and therefore identify where the CE needs to put its attention. (The "Practical Methodology" section, below, is an example of this approach.)

Quantitative Approach

A quantitative risk analysis attempts to assign monetary values to the potential losses that might occur as a result of a threat exploiting a vulnerability.

A quantitative risk analysis requires that information assets be valued by some sort of common standard (see "Sample Quantitative Risk Analysis Calculation," below). There are typically three elements that determine the value of an information asset:

1. Initial and ongoing cost of purchasing, licensing, developing, and supporting the information asset
2. Value of the information asset to the organization's operations, research, and business model viability
3. Value of the information asset established in the external marketplace and the estimated value of the intellectual property such as trade secrets, patents, or copyrights

For healthcare, such an evaluation process may be very difficult and not fully applicable. The first value should be something that can be identified but is not commonly available. While the second value is applicable, it is very difficult to determine, and can probably only be estimated.

For example, one would have to conduct an exercise to determine if certain patient information is not available in an information system, would it be available elsewhere and, if so, in a timely manner? Is it likely that the absence of the information will result in harm? Can the value of the harm be measured? Finally, value in the marketplace generally does not apply to healthcare, though if the first two values could be determined, this value being zero would not be a problem.

After information assets are valued, the rate of occurrence of threats exploiting vulnerabilities must be quantified. For example, a vulnerability may be a port left open for the information systems vendor to provide patches and troubleshoot system problems. How frequently, then, does this open port get exploited by others who do not have legitimate access? This would need to be quantified, such as once a year or once a week.

From the value of the loss and the rate of occurrence of threats exploiting vulnerabilities, annualized loss expectancy can be determined. This would be the monetary value of the risk. The result of this analysis can then be compared to what it would cost to institute security controls that would reduce the risk.

Practical Methodology to Risk Analysis, Management

While the quantitative approach may produce a seemingly straightforward answer as to when risk mitigation strategies should be employed, it is time-consuming to use and highly dependent on estimation. The principles behind the process, however, are sound and can be used as part of one's thinking process while conducting a more qualitative or ranking-oriented risk analysis.

Practical approaches to conducting and documenting a risk analysis for the HIPAA security rule may be to:

- Inventory information systems, their present security controls, and criticality of the applications and their data. Understand senior management's risk profile
- Identify threats in the environment
- Identify vulnerabilities that threats could attack
- Determine the probability that a threat could attack a vulnerability, analyze the criticality of the impact, and summarize the risk
- Determine risk mitigation strategies, implement applicable controls, and report residual risk to senior management
- Document the process
- Using information from an information system activity review, track results of controls; monitor changes in the environment, information systems, and security technology; update the risk analysis; and implement any other controls

Inventory Information Systems

A Y2K inventory or information systems review for the privacy rule may already be available. It is a good idea to track date acquired, license, location, vendor, maintenance agreements, current version, functions, and data owners or system administrators for all major hardware, operating systems, and application software. Such an inventory should encompass all information systems, whether they are located in the data center or not.

For each application, identify and describe the security features. These would include the ability to support unique user identification, access controls, emergency mode access, automatic logoff, audit controls, authentication, data integrity, encryption, and backup. Reference the policy and procedure in which each control is documented. For example, under authentication, you may describe that the application supports only up to a six alphanumeric character password (no special characters).

Part of contingency planning is assessing the relative criticality of specific applications and data in support of other contingency plan components. Criticality can be rated on a high-medium-low scale. For example, the criticality of an intensive care information system is probably high, where the criticality of an order communication system is medium (for example, orders are recorded on paper and it is feasible to call or send copies to ancillary departments).

A good way to assess criticality is to determine the amount of time the system could be down before patient care would be affected (high criticality), before operations would be significantly impacted (medium criticality), or if paper/independent computer systems could be used to load batches at a later time (low criticality).

Identify Threats

A threat is an indication or warning of trouble. In security, a threat is anything that could harm information. Most security experts consider that there are three components to a threat:

- Target—the object of a threat. HIPAA identifies confidentiality, integrity, and availability as potential targets of a threat. Many add "accountability" to the list
- Agent—the motivation or resources for carrying out the threat. Motivation is a human characteristic and could include accidental threats such as input errors, inappropriate activities such as wasting corporate resources, and intentional and illegal acts such as theft and sabotage. These may be internal or external. There are also natural acts and environmental threats that are the source of power outages, building explosions, etc.
- Event—the result of a threat. In general, there are four types of events: unauthorized access (breaching confidentiality), modification (causing a data integrity problem), denial of service (rendering data unavailable), and repudiation (the inability to identify the source and hold someone accountable for an action)

The healthcare delivery system has typically taken the approach that "this won't happen here." Unfortunately, there have been many security-related incidents that have reached newsworthy proportion, many more that have been handled without media attention, and some that may even have gone unnoticed. As part of HIPAA's risk analysis, your organization should make an honest appraisal of what might realistically happen.

Consider the location of the facility. Is the data center located over a flood plain? Are utility cables above ground where damage is more likely to occur? Some threats may be temporary. For example, is there new construction nearby? Understand the community and population that serve as a source of work force members and patients.

Some threats result from information system configuration. It may be that a best of breed environment is more prone to threats than a best of fit simply because there are fewer different people to deal with and things to learn about and remember to do.

Identify Vulnerabilities

A vulnerability is a flaw or weakness in information system security policies and procedures, design, implementation, or controls that could be accidentally triggered or intentionally exploited by one or more threats previously identified. Many healthcare organizations performed a "gap analysis" or "privacy and security assessment" as part of their privacy rule preparedness. This is a good place to start to identify security vulnerabilities in information systems. Other internal sources may be auditors' reports and risk management reports.

As privacy complaints begin to occur, review them to determine if there is a security component. Privacy and security may have a cause and effect relationship. Another approach is to simply conduct a walk-through, both literally in the physical environment and figuratively in terms of information systems. Vulnerability scanning tools can be used and penetration tests performed to identify vulnerabilities in information systems.

External sources of information about vulnerabilities include hardware and software vendor Web sites that might describe incidents others have had and provide patches or service packs to mitigate some of these. Many security associations produce online and print newsletters. Even local business groups, colleges or universities, and the police department may be good sources of information.

Summarize Risk

Typically, threats are paired with vulnerabilities, although it is not necessarily a one-to-one relationship. Many threats may exploit a single vulnerability. One threat may exploit many vulnerabilities.

One way to ensure your organization is both complete in your assessment and comprehensive in your compliance with the security rule is to document vulnerabilities and threats as they relate to each security standard. Once you have the threats and vulnerabilities documented, assign a high-medium-low rating to each threat/vulnerability pair with respect to probability and criticality. See "Rating Risk," below.

The ratings may be combined to create a numeric ranking as shown in "Risk Ranking Scale" below. This type of chart can be used to summarize the risk and prioritize remediation. Bear in mind, however, that every security rule standard must have a compliance plan. Implementation specifications may be required or addressable. The risk analysis is critical in determining how you will treat the addressable implementation specifications. Based on the risk analysis, it may be necessary to address the implementation specification as described in the rule. Alternatively, a medium to low ranking may suggest that an alternative is satisfactory.

It is possible that an implementation specification does not apply. For most hospitals and large physician offices, however, most implementation specifications will probably apply. Remember also that even though an implementation specification may not apply, the standard itself must be complied with.

Risk Mitigation

Once risk is understood, risk mitigation strategies can be developed and controls implemented. The NIST "Risk Management Guide for Information Technology Systems" lists six options for risk mitigation[7]:

- Risk assumption—the acceptance of the potential for risk. Controls may be used to lower risk, but not to the extent they could be if more resources are applied. This may be an acceptable strategy if risk is determined to be low and the cost of mitigation is high
- Risk avoidance—the act of eliminating the risk cause. Generally this means forgoing certain functions in the system or shutting the system down. This strategy is not often used, but may be necessary on a temporary basis
- Risk limitation—the implementation of controls that minimize the adverse impact of a threat exploiting a vulnerability. These controls would help deter, detect, and react to a potential threat
- Risk planning—the management of risk by prioritizing, implementing, and maintaining controls. This is essentially the process of conducting risk analysis as outlined here
- Research and acknowledgement—the acknowledgement that a vulnerability exists and the process to research appropriate controls. This should be considered a temporary strategy reserved for use during the implementation phase of the security rule, the implementation of a new information system, or when a completely new threat becomes known
- Risk transference—the selection of other options to compensate for loss, such as purchasing insurance. This generally will be used in combination with other strategies.

These options recognize that controls cannot totally eliminate risk. In general, controls may be categorized as:

- Preventive—inhibiting a threat, such as by access controls, encryption, and authentication requirements
- Deterrent—keeping the casual threat away, such as strong passwords, two-tiered authentication, and Internet use policies
- Detective—identifying and proving when a threat has occurred or is about to occur, such as audit trails, intrusion detection, and checksums
- Reactive—providing a means to respond to a threat that has occurred, such as an alarm or penetration test
- Recovery—a control that helps retrieve or recreate data or application, such as back up systems, contingency plans

In addition to what the security controls address, control strategies should include administrative, physical, and technical components. The HIPAA security rule standards themselves offer guidance on what general types of controls are required.

The security rule also references two other NIST Special Publications, "Generally Accepted Principles and Practices for Securing Information Technology Systems" (800-14) and "Underlying Technical Models for Information Technology Security" (800-33). The NIST Web site features many other helpful special publications. Many information system vendors have also posted information about what plans they have for enhancing security features.

Policy provides the overall direction for the controls. If senior management directs that controls should be preventive to the extent possible, then controls must necessarily be stronger than those where management indicates deterrent or detective controls are adequate.

Procedures will spell out the details of how specific controls will be implemented. While security policies should be known by all members of the work force, some security procedures may need to be considered sensitive information. In this case, only a

limited number of persons with a need to know should have access to procedures such as how to set passwords or the encryption methodology employed. While procedures may be sensitive and the number of persons with a need to know limited, there should always be more than one person who knows each procedure (backup) and no one person should know all procedures for all controls (separation of duties).

A final step in risk mitigation is to determine and report residual risk to senior management. Because no system can be made risk free, residual risk is that risk remaining after the implementation of new or enhanced controls. In an age of due care and ultimate responsibility for mission accomplishment, an estimate of residual risk should be made and presented to senior management. If the residual risk has not been reduced to an acceptable level, the risk analysis cycle must be repeated to identify a way of lowering the residual risk to an acceptable level.

Just as with the risk analysis itself, residual risk can be qualitatively or quantitatively described. For example, a security expert may indicate that the probability of a threat exploiting a given vulnerability is less than 10 percent. Residual risk may be described in monetary terms if a quantitative risk analysis was performed.

In the example used in "Sample Quantitative Risk Analysis Calculation," below, the estimate of asset value was based on a breach of confidentiality, a privacy complaint to the Office for Civil Rights, and a settlement, which cost the organization $500,000. The control recommended would prevent most hackers. Additionally, a security expert could offer a judgment on the likelihood that a social engineer could spoof members of the work force and get them to open the port to an unauthorized hacker.

It is conceivable, however, that the cost of the remediation could be even higher if a major lawsuit was involved. A risk manager might provide an estimate of $2 million. The likelihood of a threat still occurring with the control and the higher cost estimate would be residual risk information that senior management should understand.

Document

HIPAA requires documentation of the risk analysis and that it be retained for six years. Documentation is critical in proving that the analysis was performed. Even if you are in the "research and acknowledgment" phase, it is good practice to document exposures.

Once again, HIPAA does not specify the form of documentation a risk analysis should take. Many organizations will use some type of spreadsheet, especially if they have ranked risks or use any quantitative approach. The format in "Risk Analysis and Management Documentation," below, may be used to create a spreadsheet to document the risk analysis and conduct ongoing risk management.

Risk Management

Risk management is the act of implementing the security measures. It also entails monitoring for changes and responding with enhanced strategies. The security standards general rules also address maintenance (§164.306(e): "Security measures implemented to comply with standards and implementation specifications adopted...must be reviewed and modified as needed to continue provision of reasonable and appropriate protection of electronic protected health information."

The Information System Activity Review implementation specification under the security management process standard requires records of audit logs, access reports, and incident tracking reports. These and other internal and external documents should be periodically reviewed to determine if risk has increased. In addition, technology itself changes. Where it may have been difficult and costly in the past to institute single sign-on, new standards may make it easier to implement this measure that helps users manage their authentication process.

If specific reports do not trigger a review of risk, it may be suitable to institute specific indicators or future review dates. Federal government agencies are required by law to reassess risk to information systems every three years. This is a good benchmark from which to determine an appropriate time frame.

One measure that the final security rule does not explicitly address is configuration management. The rule's preamble explains this was eliminated as a separate standard (previously included in the proposed security rule) because it was believed to be

incorporated in other standards. Configuration management is essentially change control. Many organizations apply configuration management to information technology to manage versions of software and prioritize requests for changes to systems. A formal change control procedure should also address security. Any time a change in a system takes place, two key elements should be reviewed:

1. Are security controls in place? Were any security controls temporarily shut off to install an upgrade? Have they been reinstated? Are there default controls in the new system that should be customized to your environment?
2. Should new controls be adopted? Are changes to the system or new systems such that old controls don't work or newer controls will apply? Are there additional controls that are needed for the upgrade or new system?

Keys for Success

The NIST "Risk Management Guide for Information Technology Systems" concludes with some suggested "keys for success." In summarizing these and offering practical guidance to HIM professionals, a successful risk analysis and management program depends on people—people given the authority and assuming responsibility for complying with policy and following procedure, for awareness and reporting incidents, and for offering suggestions for mitigating risk.

The security rule contains many more administrative and physical safeguard standards than technical standards. Even as it only addresses protected health information in electronic form, it is people that make security happen.

Sample Quantitative Risk Analysis Calculation

Information Asset: Lab results from remote access reference lab

Vulnerability: Open port

Threat: Exploitation of open port by disgruntled former employees of the hospital or lab to gain unauthorized access to lab results

Calculation: Value of asset: (1) Amortized initial cost plus annual costs + (2) Cost to process an OCR complaint, fine for one violation, and risk manager's estimate of likelihood and cost of a lawsuit = $500,000

Rate of Occurrence: Since the system was installed two years ago, there has been one successful hacker who obtained test results that resulted in a breach of confidentiality, a privacy complaint to the Office for Civil Rights, and an out-of-court settlement. Rate of occurrence is one every two years, or 0.5.

Annualized Loss Expectancy: $500,000 x .5 = $250,000

Analysis: It has been determined that the only viable risk reduction strategy would be to keep the port closed at all times and have the reference lab call to request a staff member open it when the lab is ready to send information. This has been determined to cost $50,000 per year in staff time. Because the cost of risk reduction is less than the annualized loss expectancy, security experts would recommend adopting the security control.

Note: If there had been no exploitation of the vulnerability, the annualized loss expectancy would have been 0, in which case, a cost of $50,000 to introduce a control could be considered an unnecessary expense. Some security experts suggest that the potential for exploitation of a threat be estimated and used instead of actual occurrences. The potential for threat could be estimated from other healthcare providers' experiences, if known, or an estimate based on senior management's risk profile (how risk tolerant or risk averse they are). For example, senior management may indicate that in all cases, a factor of .2 should be used to estimate annualized loss expectancy, in which case the result in this example would be $100,000, still within the range where the control should be implemented.

Reprinted with permission from Margret\A Consulting, LLC

Rating Risk

Rating	Probability	Criticality
High	• You have experienced an incident • Controls are not very effective	• Results include human death or serious injury • Inability to recover critical data or high cost of recovery • Major lawsuit • Loss of licensure or accreditation
Medium	• You have been alerted to threat • Controls may impede threat	• Results include human injury/harm • Complaint to federal government • Significant cost of recovery • Minor lawsuit • Public relations issue
Low	• No one in community has experienced threat • Controls will greatly deter or prevent success of a threat	• Complaint • Loss of productivity • Nuisance • Embarrassment

Reprinted with permission from Margret\A Consulting, LLC

Risk Ranking Scale

Rating	Criticality		
High	3	6	9
Medium	2	4	6
Low	1	2	3
	Low	Medium	High

Risk Analysis and Management Documentation

Standard	Vulnerability	Threat	Probability	Criticality	Risk Score	Control	Residual Risk	Review Date
ß164.308 (a)(1)(ii) (A) Risk analysis	Only gap analysis performed to date	Existence of threats not documented	M	H	6	Risk analysis	Not all threats may be identified	April 20, 2004

Prepared by

Margret Amatayakul, RHIA, CHPS, FHIMSS

Notes

1. “Health Insurance Reform: Security Standards; Final Rule.” 45 CFR parts 160, 162, and 164. *Federal Register* 68, no. 34, page 8377 (February 20, 2003). Available at http://aspe.hhs.gov/admnsimp/.
2. *Ibid.*, page 8364.
3. *Ibid.*, page 8377.
4. *Ibid.*, page 8343.
5. NIST. Special Publication 800-30, “Risk Management Guide for Information Technology Systems.” Chapters 2 and 3. For more information, visit www.niap.nist.gov.
6. *Ibid.*, page 6.
7. *Ibid.*, page 27.

Article citation:
Amatayakul, Margret. "Security Risk Analysis and Management: an Overview (AHIMA Practice Brief)." *Journal of AHIMA* 74, no.9 (October 2003): 72A-G.