CHAPTER

Risk Assessment Techniques

10

INFORMATION IN THIS CHAPTER

- Operational Assessments
- Project-Based Assessments
- Third-Party Assessments

INTRODUCTION

Once you have a risk model and a few assessments under your belt, you will want to start thinking strategically about how to manage the regular operational, project, and third-party assessments that will occupy most of your time as a risk manager or analyst. This can quickly become an overwhelming task if not approached strategically, making the best use of the tools and resources that are available. You will want to have a single risk model for the organization, but the actual assessment techniques and methods will need to vary based on the scope of the assessment. An assessment of risk during an incident investigation, for example, must be more streamlined than an architectural risk assessment of a new software application in development.

OPERATIONAL ASSESSMENTS

Do you think that you would use the exact same techniques to perform a risk assessment on a new application or system in development as you would use to assess an entire company during an acquisition? The answer is that you wouldn't. So far, we have established risk models and frameworks, which will be the foundation for any assessment, but how you go about performing that assessment will vary based on the size and nature of the target. It can be helpful to start thinking about categories of assessments, beginning with the distinction between operational assessments, meaning those ongoing day-to-day assessments that are occurring all year long, and project-based assessments which have a finite duration. The operational assessments will encompass regular assessments of emerging threats, newly announced vulnerabilities, and discovered standard violations, just to name a few. Operational assessments should not be confused with assessments of risks in the operations domain. In this context, operational describes the format

Security Risk Management. DOI: 10.1016/B978-1-59749-615-5.00010-4

of the assessment, indicating that these are ongoing and revolving assessments with no clear endpoint, as opposed to assessments of projects that have set completion dates. In contrast, an assessment of the operations domain would define the scope of the assessment, which would focus on threats to operations continuity. We are focusing on the former for the purposes of this discussion.

Some examples of operational risk assessment tasks in the information security space include the following:

- Threat analysis
- Vulnerability scanning
- Patch remediation
- Penetration Testing
- Incident prioritization
- Exception processing
- Compliance to standards reviews
- Certification and accreditation (C&A)
- Auditing (internal or external)
- Responses to client due diligence evaluations
- Vendor on-site reviews
- Regulatory gap analysis

As you can see, this list is rather diverse, and even so, it doesn't even begin to cover all the various tasks for which a security risk management team might be responsible. It just wouldn't be practical to use the exact same approach and techniques for each of these tasks, but fortunately, the fundamentals stay the same. It is really just the tools and format of the assessment that change with the type of task. For example, a vulnerability scan of your Internet presence is going to require a technical tool or service to perform security scanning of vulnerabilities, but an on-site review of a service provider's physical security controls is going to require a body with a clipboard and a list of required controls. Likewise, you aren't going to require an on-site physical assessment of Dell's facility just because they provide your server hardware, but you would want to perform that on-site assessment of an offshore development center that provides 80% of the code for your products. When you are establishing your risk management program, start by thinking about the different levels of resources that you will be assessing and map out which methodology will be most efficient for each.

Operational Techniques

For all those potential operational assessments, your options really come down to just a few assessment formats:

- Questionnaire
- Interview
- Passive testing

- Active testing
- Review of third-party assessment
- Acceptance of a certification

When it comes to internal or third-party assessments, you should consider mapping the depth and intrusiveness of the assessment technique to the risk sensitivity of the service being provided. For example, a review of an independent assessment report or a passive test, such as conducting a Google search for information about your organization, will usually be nonintrusive, requiring mostly only your own team's resources. For those resources that have lower risk sensitivities or have already been reviewed in the past without any significant findings, you may want to consider these approaches to minimize your impact on staff from other business units.

Questionnaires and Interviews

The first two techniques are questionnaires and interviews, and we will address them together since, ultimately, a questionnaire is just a passive version of an interview. Choosing which is appropriate can often be difficult and it may come down to trial and error to determine which one your organization responds to better, but hopefully, these guidelines will give you a good place to start. First, the benefit of an interview style assessment versus a questionnaire is that a skilled assessor can use the responses to a static question to guide their follow-up questions and direct which additional questions they ask. For instance, if you are assessing the IT environment and you have a series of questions about password controls (length, complexity, change history, expiration, initial distribution, reset procedures, and so on), but the system in question uses digital certificates or cryptographic keys instead, you can skip all the remaining password questions and drill into the key management questions on the fly. To do this with a questionnaire, you either need to program some logic into an online questionnaire or you will be doing a lot of back and forth follow-up questions about why they selected "N/A" for all your password questions.

Especially, if you are doing an internal assessment, you would be surprised how many additional risks you can uncover just by getting several people in a room at once and listening to them disagree about how something actually works. The manager will give you one answer, the engineer will correct him, and the junior engineer who recently joined the team will say "nobody told me that was the procedure." Of course, the above scenario assumes that some level of trust has already been established, that the culture supports healthy disagreement in public, and that your assessor understands the power of just listening. A side benefit of the interview technique can often be increased awareness among the team being assessed about what is expected from a security perspective and, as a result, bad practices can often be corrected right then. In contrast to that situation is the defensive interviewee or the subject who is actively offended that anyone would dare question their practices. If you suspect that might be the case, then a questionnaire might be the more effective way to go.

No matter how long you spend crafting the "perfect" questionnaire, you will always have questions that are misunderstood. If the question isn't clear, you will probably experience one of the following responses from the person answering the questionnaire (in order of likelihood, from most to least likely):

1. Skip the question altogether
2. Select "N/A" if it is an option
3. Give up on the questionnaire entirely and not finish it
4. Answer the question with a "No" just to be safe
5. Ask for clarification

You may wish that response 5 was more common, but with so many pulls on resources' time, you are probably going to have to hunt down the responder to find out that there was a question they didn't understand. You can minimize this situation by trying to provide organization-specific examples along with each question. A targeted example can go a long way toward clarifying the intent of the question. Of course, when conducting an interview, you can address any confusion immediately, which minimizes the time lost and the frustration experienced by both sides.

As a general rule, using an interview style is going to give you the richest and most accurate information in the shortest amount of time, assuming you can get the right people in a room all at once. It may seem onerous to schedule all these interviews and coordinate resources, but it gets you exposure to many critical functions in the organization and will be your quickest option. The challenge is that interviews don't scale well for large organizations, so you will need to prioritize where you use a questionnaire versus an interview. One approach is to use an interview for the first assessment and a questionnaire for each subsequent assessment for that same resource. That way, you get a detailed risk assessment and understanding of the resource up front, but can scale back the resource effort over time. Another approach is to send out a questionnaire and schedule an in person meeting with everyone involved to review the answers and discuss any follow-up questions. With this approach, you leverage the benefits of both assessment formats.

Active and Passive Testing

Questionnaires and interviews might work well for identifying policy violations or process weaknesses, but to really evaluate the technical vulnerabilities in your environment, you will need to perform some sort of security testing. Although passive testing sounds harmless, beware that the definition of passive is not always consistent across the field. There are definitely gray areas to be aware of; any testing should require appropriate senior management approval. Most security scanners or vulnerability scanners are tools with large databases of known attacks and weaknesses and will scan the environment for signs of vulnerabilities or compromises. These tools will also typically have the ability to identify missing patches, configuration mistakes, or denial-of-service weaknesses.

Security scanning tools are very common. Many will focus on general operating system and commercial application vulnerabilities, but others specialize in mapping environments or testing Web applications for weaknesses. Most will only look for signs of a weakness, while others also include the option to validate a vulnerability by actually exploiting it. Any tool that will actually verify a weakness by executing the exploit would be considered a penetration testing tool, not just a scanner. There are many open source and commercial scanners available. A few of the most common ones are as follows:

- Nessus (free and commercial versions available)
- NMap (free)
- ISS
- Retina
- Nexpose
- Foundscan
- Qualys
- Core Impact
- AppScan
- WebInspect

This list doesn't even come close to being inclusive, especially as you start to look at specialized scanners for targets like wireless networks and Web applications. A great list of the top 100 network security tools is available on Gordon Lyon's SecTools site [1], and many of these tools are security scanners of some kind. Gordon is the author of the NMap scanner, so he knows a little something about the topic.

The scope of an active or passive test can range greatly depending on your organization's particular concerns. For example, the following are all typical types of assessments:

- Enterprise vulnerability assessment (active)
- Penetration testing analysis (active)
- Wireless security assessment (active)
- Blackbox application testing (active)
- Malicious threat assessment (passive)
- Internet reconnaissance (passive)
- Application code security review (passive)

Most of these should have an obvious scope; however, *malicious threat assessment* and *Internet reconnaissance* both likely need some further explanation. Typically, a *malicious threat assessment* would involve putting a passive security device at key network aggregation points to review traffic for potential malicious activity or policy violations. This is sometimes accomplished by temporarily putting a specialized Network Intrusion Detection System (NIDS) device, or an anomalous network activity monitoring device like the Riverbed Cascade (formerly known as Mazu) analyzer, on the network, and reviewing the alarms

that are triggered. This is a passive test because at no point is there any chance that the normal operations of the network can be impacted. Signatures and anomaly detection techniques aren't perfect, so it may be useful to conduct one of these tests every so often, even if you already have intrusion detection systems (IDS) deployed in your environment. Just having an analyst look at your network traffic for a week without the prejudices of what is expected or suspicious can often uncover unknown issues.

WARNING

No matter what kind of testing is proposed and how much the tester assures you that there will be zero impact to your environment, be very cautious. Even the deployment of a passive monitoring device on your network could impact operations if, for example, the device is accidently assigned an IP address that is already being used by another critical server. This may sound implausible, but be assured—it really happens! Better to be cautious and run installs of even passive monitoring devices through proper change management processes.

An *Internet reconnaissance* test should be focused on assessing the organization's profile based on what information is publicly available on the Internet. Domain registries, the organization's financial statements, career postings, and vendor case studies are all sources of information about an organization that could be used by an attacker. Google has actually become a primary tool for would-be attackers to profile an organization looking for weaknesses that can be exploited by technical means or through social engineering. Any organization needs to have some level of public presence, a point that is emphasized by the introduction of the White House as an active participant on Facebook during the Obama administration. The point of this type of testing is to have someone with the knowledge of typical data mining techniques look at the organization's profile from an Internet perspective and identify unnecessary information risks. Like other passive testing methods, this assessment presents no risk of an operational disruption to the organization.

Most active testing will involve either a tool or a person performing functions against a resource to look for known responses, which indicate that a vulnerability is present. For example, an active scan of your environment would look for known vulnerabilities and improper configurations that could allow an attacker unauthorized access to a resource. It is always recommended that you scan your environment both internally and externally so that you get an idea of what would be visible to any outside attackers as well as potentially malicious insiders. It is a good idea to publish a formal schedule for scanning and to communicate this to resource owners and administrators. You may need to do your scanning during off-hours or maintenance windows to avoid affecting a production service. After all, no matter how much time you put into tuning your scanner, you can't guarantee zero impact to the environment being scanned, and resource administrators need to be prepared to respond if needed to a disruption.

One focus of security testing needs to be to validate that current controls are behaving as expected. It isn't enough to just implement a set of controls; you need to evaluate those controls to ensure they are really reducing your risk exposure to the level you expect. Controls also require constant tuning and adjustment, especially with the growing sophistication and persistence of attackers, and you will need to be constantly monitoring each layer of controls to see which attacks are getting through. If you think that your firewall is locked down, run a port scan to verify. If you are relying on your anti-virus software to catch the latest threats, introduce a few sample pieces of malware into an isolated and controlled environment to see the detection rate (virtualization with no network connectivity can be a great test bed). If you think that peer review of application code is catching the violations of coding standards, have a security architect review a random sampling of code to validate. As they say, trust but verify.

In addition to regular scanning and other internal assessments, it is crucial to have outside experts come in periodically to assess different parts of the security program by performing penetration testing on the network or Web application, or by trying to bypass physical controls like gaining access to a secured area without a badge. This will help you to identify weak areas that need more attention and can also help you validate the threat vectors that you have assessed as most likely.

Third-Party Reviews and Certifications

When working with vendors and service providers, you are going to need to rely on other means of assessing the security posture of the third party. Most service providers aren't going to let you show up at their offices with a security scanner and just let you go nuts on their environment (at least we hope they won't!). Thus begins the negotiation of best evidence. You might think of this as a similar dilemma to what you would see in court. Direct evidence may not always be available, so you may need to rely on alternatives like maybe an expert witness. The same is often true when assessing a third-party provider—you may not be allowed to walk through their Security Operations Center (SOC) or run your own penetration test against their Internet-facing systems, but they should provide you some indication that they have had an independent third-party assessor perform these tests and that any high-risk findings are being addressed appropriately. The debate about the appropriate level of detail to require will be discussed in depth later in this chapter, but suffice to say for now that you likely shouldn't expect a copy of a penetration report, but it might be reasonable to request an executive summary. After all, the provider also has to manage the risks inherent in distributing active exploit details.

If report summaries from independent assessors are not available, the next best thing would be a certification that demonstrates a certain level of security posture and program maturity. For example, you might recognize an ISO 27001 or SAS70 Type II certification as being sufficient proof of robust security controls for the organization. Eventually, the industry will need to develop a certification that covers all the areas of review in the 800 to 3,000 question evaluations that

some customers are requiring their providers to complete, but as a field, we aren't there yet. The SAS70 certification, for example, can be a fantastic evaluation of security controls, but the scope will vary between organizations depending on what they chose to include in the review and the level of detail in the report. This makes the certification hard for risk managers to use as a consistent indicator of excellence.

Baseline Reviews

In terms of operational risk assessments, another important focus is Certification and Accreditation (C&A). For many business professionals, these terms may not be meaningful, but don't worry: like with the term *information assurance*, you will most often see these terms in the context of the US federal government. Although the terminology isn't popular in private industry yet, the function actually is already in use. On the most basic level, C&A tasks require establishing a security baseline for each system in your environment, ensuring any new deployments are compliant with the baseline, monitoring the configuration of the system over time to be sure it doesn't deviate from the baseline, and documenting any areas where the system can't comply with the baseline. In essence, a C&A process is meant to formalize the standards for configuring a system securely and force an explicit review of those controls and authorization decision to allow it to operate in an environment.

Certification and accreditation are really both subsets of an overall *information security risk management* program. Risk management is the overall program for identifying weaknesses, threats to those weaknesses, and assessing the impact to the organization that might result from an exploitation of those weaknesses. Certification is the process of evaluating whether the system/application meets the minimum standards that have been established, and *accreditation* is the management decision process to determine if any deviations from standards are acceptable. When you think about this in basic terms, it essentially equates to a risk assessment followed by a risk decision. In the US federal government, there are very explicit job roles and positions involved in this process; however, most corporations use a combination of the resource owner or operator and a representative from the security team to negotiate these details.

There are two contexts in which the term "baseline" is used for Information Security. The first is referring to a point in time snapshot of the current state of the environment as a comparison point. The second is the minimum set of required configuration settings or controls to meet a desired level of security. In this chapter, we are using the latter definition—just think of it as a secure configuration template.

There are many activities required to make a C&A process run smoothly, and many of these tasks will be performed by the resource administrators or operations teams, with oversight from the Information Security team. As part of the change management process, the postimplementation steps of updating documentation such as network diagrams, server build documents, software hash libraries,

standard build images, and so on should be performed. A good practice is to create a hash library of known good software in your environment; that way, when there is an investigation of a system compromise, you can easily identify software and configuration files that have not been tampered with because they match the unique hash you created in advance.

Many organizations also run regular (as often as nightly) scans of server configuration files to ensure they still meet the baseline, and if any deviations are found, they get escalated to management and the security team to investigate the cause. When a deviation from the baseline is required due to technical constraints or for specific business purpose, the justification, risk evaluation, and approval needs to be documented and processed like any other risk acceptance. This assessment needs to happen before the system/application "goes live" or is released and regularly until it is decommissioned. It is important to ensure that this requirement is communicated to all project managers and stakeholders so that they can account for this time upfront when they create a project schedule. You will also need to establish who has the authority to keep the system/application from going live or being released if there is a serious security issue. You may hear this authority referred to as the "red lever." This is the person who the organization has established as having the authority to stop a system from going into production or to shut down an existing system if the exposure warrants it. Accreditation is not a permanent state; the security of any system/application needs to be re-evaluated periodically, usually on a set schedule, the frequency of which should depend on the sensitivity of the resource. The NIST Special Publication 800-37 Revision 1 [2] is a great reference for anyone who is involved in C&A work. It has evolved in this revision from a rather static and inflexible process into a risk-focused lifecycle methodology.

Assessment Approaches for Different Sized Scopes

When you are faced with assessing a very large environment, "random sampling" should be the first words that come to mind. It may be feasible to perform full port and vulnerability scans on 20 systems in a reasonable amount of time without putting a dangerous load on the systems or the network, but think about the logistics if you needed to assess 2,000 systems. At that point, is there really any value in documenting the same weaknesses across all 2,000 systems? You don't need a sample size that big to establish a pattern. Especially if you are in a consultant role, you will want to very carefully consider what scope of assessment would be a productive use of time and resources. Remember that whatever you test, you will need to document and report on. Because of these considerations, and just like auditors have been doing forever, random sampling is the best approach.

Similarly, there are often debates about whether automated penetration testing is sufficient for a thorough assessment, as opposed to having a highly skilled ethical tester hacking away at the application or system manually. Clearly, the latter option is preferred, but it is also typically not possible as the only method of

testing on large-scale assessments. If you are looking to assess a specific function in an application that uses a proprietary protocol, then maybe a purely manual penetration test is the right solution, but for large-scale assessments, any tester is going to use a hybrid of some level of automation along with manual testing.

PROJECT-BASED ASSESSMENTS

Chapters 11 and 12 will cover daily risk assessment activities that continue on a constant cycle, but for now, let's first look at how best to approach an assessment with a defined endpoint based on a single project. The three most common projects that will require a risk assessment are as follows. Each requires a slightly different approach and has its own challenges.

- Software development
- Software/technology acquisition
- Selection of third-party service provider

The scope of an assessment can vary greatly, from a new product enhancement to the acquisition of another company. However, the process and deliverables are going to be the same, even if the subject matter varies. The important distinction here is that this is a point-in-time assessment and not an on-going process like operational assessment activities, which we have previously discussed. Because of this, it is necessary to have a set project timeline and clear deliverables to guide the assessment.

Risk Assessments in the Project Lifecycle

Generally, the motivation for this type of risk assessment will be to demonstrate due diligence and assess the level of risk being undertaken by the project. These assessments need to be performed as early in the project's lifecycle as possible so that it can be properly influenced by the results of the assessments from the beginning. Otherwise, time and effort may be lost if the team is allowed to go too far down a flawed path. A security risk assessment can be performed by just about anyone involved in the project team if given the proper guidelines, and occasionally, the project may require an outside party to guide the assessment. Your organization's culture will strongly influence who should lead each assessment, but generally the responsibility will fall on the Information Security team.

The output of this assessment will include the identification of risks, threats, and general concerns from the team and, ultimately, recommendations for controls to mitigate those threats. The analysis and recommendations would then generally be presented to senior management or other project stakeholders to make the final decisions. You should notice that this is no different than any other assessment methodology that has been introduced so far, so what really distinguishes the project-based assessment is that it is time-boxed and is designed to be a point-in-time evaluation.

The FRAAP Approach

If you are interested in a structured approach to an accelerated assessment, Thomas Peltier has coined the term Facilitated Risk Analysis and Assessment Process (FRAAP) [3] to describe his approach to managing a risk assessment of a project in a short timeframe. Using this streamlined approach, you can cut down the time it takes to gather risk data and produce recommendations, while still getting the appropriate Subject Matter Experts (SMEs) involved. The goal is to conduct the assessment in a matter of 4 to 8 hours and then produce the recommendations within a few days of the assessment session. This can really help to keep projects on track and minimize the time requirements on the SMEs. Within this model, it is especially important to define a structured agenda and strict roles for each participant. In doing so, you can avoid risk discussions that can drag on and drift far from the focus area. The role of the Information Security representative is to facilitate the discussions rather than dictate the direction they take. Depending on the topic, there might be several business units or departments represented in the session. Some of them are listed here as follows:

- Functional owner
- Business analyst
- System engineer
- Database administrator
- Network administrator
- System programmer
- Application programmer
- Functional manager
- Information security
- Legal
- Human resources

If you are serving as the facilitator for the meeting, then it is best to have someone else on your team attend to represent Information Security. This way you can focus on the facilitator's responsibilities and not be perceived as pushing your own agenda. The idea behind the FRAAP format is to run a session that encourages the participants to raise issues and identify risks, without spiraling out of control with side discussions and tangents. Once the risks have been identified, the team analyzes the impacts and agrees on the likely consequences of those risks. Then, each risk is rated in terms of the priority to the organization.

For the most part, this analysis relies on the expertise and knowledge of the people in the room, including representatives from the security team, but it can also be influenced by other data about the resources or observed trends in the industry that were gathered prior to the assessment session. Activities like brainstorming serve an important role in Peltier's FRAAP approach, but there is also enough structure to the assessment format that it should be able to keep the session productive.

Prep Work

Before you lock everyone in a room for 8 hours talking about risks, you will need to do some preparation work. Set up a 1-hour pre-session meeting with the primary stakeholder, project lead, and session facilitator to discuss the goals, agenda, and format for the session. Keep this meeting short and focused. Peltier recommends five deliverables for this meeting:

- Draft a scope statement for the initiative and the assessment.
- Obtain visual diagrams of any resource components, inter-dependencies, or information flows.
- Select the team members for the actual assessment session.
- Decide on meeting logistics, such as location, timing, supplies, food, and so on.
- Agree on definitions of any controversial terms such as the following:
  - Confidentiality, integrity, availability, accountability
  - Risk
  - Threat
  - Vulnerability
  - Impact
  - Control

Having this defined up front and published to the assessment team will avoid wasting time at the beginning of the session trying to get everyone on the same page.

The assessment session itself should last between 4 and 8 hours, depending on the size of the project and shouldn't include more than 15 participants. If you can get everyone into a room off-site or at least away and disconnected from everyday distractions, then the sessions will be far more efficient. The last thing you want to see is everyone sitting around the table answering e-mails on their laptops or BlackBerries. Expectations need to be set early that this time will be dedicated to the project and the assessment activity.

The facilitator will want to come to the assessment session prepared with materials, such as flipcharts and markers, printed copies of the terminology definitions, a clear scope statement, and any visual diagrams that might be appropriate. It is recommended to distribute the materials from the pre-session meeting in advance. This gives the participants the opportunity to review them in advance, gather any information that they might need for the meeting, and also identify if they are not the right resource to be involved in the assessment. You should, however, also assume that the majority of people will not review the materials in advance, so plan to spend a few minutes summarizing the scope at the beginning of the meeting.

Running the Session

The assessment session itself should only last between 4 and 8 hours. You will have to consider your audience, scope of the assessment, and the culture of your organization when choosing the length of the session. Some assessments may be hard to

complete thoroughly in just 4 hours, but you also have to account for people's attention span and other draws on their time. Scheduling any large group can be difficult, so the shorter the session, the better chance you have of getting everyone together. The session itself should have three deliverable goals. There is no one single way to capture this information, so experiment with a few approaches and choose the format that works best for you.

1. Identify the risks
2. Prioritize the risks
3. Identify controls to mitigate the top priority risks

One way to start off the session is to go around the room and ask each participant to identify any risks associated with confidentiality. Set a maximum time for this exercise (say 3 minutes per person) and capture all the ideas, then go around the room again and spend the same amount of time listing all the integrity risks. Repeat this process again for availability and accountability to create a comprehensive list of risks. Alternatively, you could begin by going around the room and asking each participant to list one issue or concern that they have with the project. When you're facilitating a session, keep in the mind the usual brainstorming tips, such as the following:

- Remain neutral at all times
- Don't judge or dismiss any ideas
- Ensure that all ideas are captured
- Solicit input from everyone
- Only let one person speak at a time
- Don't let any one person dominate the conversation

To keep the session moving forward, the facilitator needs to be very strict about following these general brainstorming guidelines. Especially for security professionals, it can be hard to stay in character as the facilitator and not comment on the issues or ideas being raised as you would when wearing your security hat, but this separation of roles is important to the success of the session. It can be difficult to find the balance between allowing participants to be creative and not letting any one personality dominate the discussion. A good way to avoid this is by cutting people off after 3 minutes or so. Otherwise, you may find the session spending too much time on a single issue and missing others. Be sure to have someone tasked with recording all the ideas and issues being raised, and defer those that are out of scope for this project. Finally, be sure to manage the group so that you only have one conversation going at a time; if the debate gets heated, you may need to mediate to keep the conversation productive and above the line.

Next, look at each identified risk and analyze the severity of each. Then, take each of those risks and rate it based on the likelihood of occurrence. Following this, you will want to prioritize the risks based on their risk rating and focus the rest of the session on the higher exposure items. The assessment and analysis will

be captured in worksheets, similar to the worksheets provided as a part of OCTAVE Allegro, which we will explore in Chapter 11.

Sample Worksheets

Having a structured assessment approach is essential to the viability of the FRAAP approach; so, this section provides several worksheets that can be used to capture the artifacts of each step of the FRAAP session. Keep in mind that each worksheet has been slightly adapted from the typical FRAAP worksheet to meet the risk model used in this book, but the general concepts remain the same.

The first step in the session is to start identifying concerns or risks, assign them a risk type (C-I-A-A), and identify the resource affected by the risk. You can see an example of this in Figure 10.1.

Once you have completed the brainstorming, review the list of risks identified and eliminate any duplicates. You should now have a list of categorized risks with the associated resources identified. Next, on the *resource sensitivity profile* worksheet (Figure 10.2), you will start out by listing each resource from the first worksheet.

Once you have listed all of the resources, include a very short description of that asset's sensitivity or importance to the organization. Use this description to guide your rating of the resource's *confidentiality*, *integrity*, *availability*, and *accountability* sensitivities using the same Low-Moderate-High scale from our

Risk description list			
#	**Risk type**	**Risk description (vulnerability and consequences)**	**Resource**
0	Confidentiality	Sensitive account information is discarded in the regular trash, which could lead to disclosure of customer financial accounts to unauthorized internal or external parties. Disclosure of this data violates several state privacy laws.	Paper statements
1	_____	_____ _____	_____

FIGURE 10.1

Risk description list worksheet.

Resource sensitivity profile							
#	**Resource impacted**	**Sensitivity description**	**Confid**	**Integ**	**Avail**	**Acct**	**Overall**
0	Paper copies of client account statements	Client account statements include the client's name, financial account number, address, and current balance. This information is protected by several regulations and privacy laws. Disclosure could lead to financial fraud and liability for the organization, or legal penalties.	High	Low	Low	Moderate	High
1	_____	_____ _____	_____	_____	_____	_____	_____

FIGURE 10.2

Resource sensitivity profile worksheet.

Risk exposure rating							
#	Brief vuln desc.	Threat category	Threat activity	Like	Sev	Sens	Exposure
0	Account information in trash	External targeted attack	A criminal could pull a few client's sensitive financial account information out of the dumpster behind the office and use it for fraudulent purposes.	High	Moderate	High	High
1	Account information in trash	Internal abuse	An employee could pull sensitive financial account information for all clients out of the trash cans in the office and use it for fraudulent purposes.	Moderate	High	High	High
2	_____	_____	_____ _____	_____	_____	____	_____

FIGURE 10.3

Risk exposure rating worksheet.

earlier *security risk profile*. Finally, determine the overall sensitivity for the resource based on highest of the individual C-I-A-A values.

At this point, you have identified the risks and their associated resources and rated the sensitivity to risk for each resource. Next, you will need to break each risk into its threat and vulnerability components, as shown in Figure 10.3.

Notice in the example in Figure 10.3 that one initial risk has been separated into two different combinations of threat/vulnerability pairs with slightly different risk ratings. This illustrates how the combinations of threats and vulnerabilities can result in different risk exposures depending on the threat category. The threat categories being used are as follows:

- Natural disaster
- Infrastructure failures
- Internal abuse
- Accidents
- External targeted attacks
- External mass attacks

In this worksheet, the likelihood and severity of the threat/vulnerability pair is combined with the sensitivity of the resource from the previous worksheet to derive the final exposure rating.

You can use the qualitative mapping table from Chapter 6 to derive the exposure value from the likelihood, severity, and sensitivity ratings.

Once the risks have been captured and rated, you will need to identify the controls that will mitigate them. You can start with a list from one of the numerous industry resources available (for example, ISO, NIST, NSA) or you can build your own custom list. Often, organizations will publish a list of approved or existing controls and technologies. This can help to reduce the complexity of the environment and increase the reusability of previous investments.

A good place to start is with the *20 Critical Security Controls for Effective Cyber Defense* [4]. These Top 20 Controls were agreed upon by a consortium US government representatives, which included the National Security Agency (NSA), US Computer Emergency Readiness Team (US CERT), the Department of Defense Joint Task Force-Global Network Operations (DoD JTF-GNO), the Department of Energy Nuclear Laboratories, the Department of State, and the Department of Defense Cyber Crime Center, plus the top commercial forensics experts and penetration testers that serve the banking and critical infrastructure communities.

Use the *mitigating controls list* worksheet shown in Figure 10.4 to capture the mitigating controls that could be implemented to address the risks identified above, including the control type of *preventative*, *detective*, or *responsive*. You will use this worksheet to map each risk to the control that will adequately mitigate it. When choosing controls, follow these guidelines:

- Identify controls that address multiple risks
- Focus on cost-effective solutions
- The total cost of the control should be proportional to the value of the asset

In many cases, multiple controls will be needed to properly mitigate a single risk. Likewise, a single control may mitigate several risks. In the original FRAAP worksheets, there are a few interim worksheets to help illustrate the effects of mapping the controls to the risks and the risks to the controls. This can help you see the controls that will give you the most bang for your buck. For the sake of simplicity, those worksheets have been eliminated here and have been replaced with the single *action plan* worksheet in Figure 10.5. The last worksheet (Figure 10.4) was a simple list of each risk and the controls that could be used to mitigate it in the order the risks were identified. The *action plan* worksheet should summarize all the information that you have gathered so far for each of the priority items. These should be listed in order of importance.

You want to start by focusing on the risks with the highest ratings because they require the most immediate attention. The moderate risks will need attention soon and the low risks can be dealt with when time and resources are available. You may also focus on prioritizing the controls that mitigate the most risks. When you are recommending actions, remember to think about the time and resources

Mitigating controls list			
#	Brief risk desc	Control type	Control description
X	Insider stealing paper statements	Preventative	Paper cross-cut shredder in all the mail rooms
Y	Insider stealing paper statements	Preventative	Data classification and handling policy, requiring the use of a shredder for all sensitive documents
1	_____	_____	_____

FIGURE 10.4

Mitigating controls list worksheet.

Action plan								
#	Brief risk desc	Risk type	Rating	Control	Priority	Owner action	By whom	When
1	Insider stealing paper statements	CONF, ACCT	High	X, Y	6	Buy a shredder and install in convenient location, and publish a handling policy	Administrative staff	4/30/07
2								

FIGURE 10.5

Action plan worksheet.

that will be required to execute on the plan. There is no value in listing action items that aren't practical.

Be sure to identify who is responsible for each item and include a deadline. As you are considering mitigating controls, always keep in mind that accepting a risk as-is may be an option as well.

Reporting

After the completion of the assessment session, your goal should be to have the report ready within 4 to 6 days. Because a template is being used to gather the information, it will be easier to compile into an assessment and recommendations report. The report is generally written by the session facilitator. Finally, a postsession meeting or meetings should be held with the stakeholders to present the report. This may be done in two or more meetings: one for an executive level overview and one to dive into more detail about each issue.

The FRAAP approach is just one technique for adding structure to your project-based risk assessments. Its value is in the defined agenda, roles, and worksheets for capturing, rating, and mitigating risks. As was illustrated here, the model is flexible enough to allow you to expand on the worksheets and risk scales themselves over time to incorporate them into a different risk model. If you expect to perform any project-based assessments, it is highly recommended that you read Peltier's book *Information Security Risk Analysis*.

THIRD-PARTY ASSESSMENTS

Earlier in this chapter, we started to lay out the challenges that quickly present themselves when dealing with third-party assessments. Almost every organization these days is experiencing this struggle from both sides of the relationship, as the client performing a due diligence evaluation and as the provider answering client or partner queries. This process can quickly become a huge drain on your resources if not managed properly. Let's pull back the curtain and look at the inner workings of this process from behind the scenes.

As noted earlier, the first issues you will encounter are the lack of an industry standard format for vendor risk assessment questionnaires and the lack of a universally accepted certification as an alternative to individual evaluations from clients. What makes it even worse is that even among clients who are looking for the same general information, each questionnaire will word the questions just differently enough that you can't even reuse your answers. So, you will need to have staff who are capable of crafting appropriate answers to these questions based on their knowledge of your security controls without giving away too much information, all the while responding in as positive a manner as possible so as to discourage follow-up questions. People with this skill are not easy to come by. To make matters even worse, you may get multiple questionnaires from different parts of the same company, and you can't even count on the questionnaires staying the same from year to year from a single client. As their programs grow and mature, they change the focus of their questions so that you basically end up having to start from scratch each year.

Industry Standard Assessments

Hopefully, you didn't just read that long list of issues and give up because solutions to these problems are available. Now that we have adequately framed out the context for the challenges, let's talk about these solutions. In the short-term, the solution that can have the largest positive impact would be creating a standardized set of vendor due diligence questions in a common format to eliminate the need for so many customized responses. This will greatly speed up the request turnaround time and allow service providers the ability to provide the most accurate answers possible. With so many ad hoc requests coming in, it can be challenging to always find the right SME to provide a definitive response and you can be sure that some clients ask some very obscure questions.

The good news is that there is a standard questionnaire emerging out of the financial services industry that could meet this need if adoption of it continues to expand quickly outside of this industry. Out of the BITS Financial Services Roundtable, there emerged a Standardized Information Gathering (SIG) questionnaire [5], which is aligned with the Federal Financial Institutions Examination Council (FFIEC) guidelines, and the Agreed Upon Procedures (AUP) report, which can be provided as a substitute for individual client tests of stated security procedures. The value of these tools has been proven by several large and small organizations; however, its usefulness all hinges on universal adoption of this format between businesses and their providers, and between businesses and their clients/partners. Version 6 of the SIG was released in 2010, but you will likely see organizations using a variety of versions from 3 to 5 as they transition to the newer version, which promises to have streamlined the number of questions down from several thousand to a more manageable number. The breadth of topics covered include the typical security operations

and policy questions that you might expect, as well as sections ranging from physical security and business continuity to privacy. As a standard, the SIG is far from universal, but adoption is growing at a fast pace, to the point where some of the leading GRC tools have integrated it into their software offerings.

Levels of Assessment

This section certainly isn't meant to be an advertisement for the SIG, but there are several features of the implementation that make it worth highlighting. The first is (as of version 5) the flexibility to roll out three levels of detail, a level 1, a level 2, and a detailed version. Version 5's level 1 contains around 100 questions, which lends itself really well to a client due diligence request during the early stages of a vendor review when a Non-Disclosure Agreement (NDA) may not have been established yet and sales needs a really quick turnaround. Using the level 1 questionnaire, you can quickly identify if there are any red flags or show-stoppers that would cause you to reject the vendor as a candidate, without the vendor having to give away too much about their controls. The *version 5, level 2* questionnaire is more detailed, with closer to 400 questions, so this might be more appropriate for later in the contract negotiations, either after the contract has been signed or at least when an NDA is in place.

Another way to make use of the levels in the SIG is when you are performing due diligence reviews of your own vendors and service providers. It allows you to base the level of SIG required on the sensitivity level of the service being provided. One possible schedule for assessments is shown in Table 10.1.

This particular schedule might assume that there was no high-exposure risk found during the review in the first year. You could set a threshold for the number or level of the findings to determine the frequency and depth of the assessments performed. Another consideration to keep in mind is that vendors in your category of low sensitivity might not require a full formal review at all. Vendors who, for example, just provide you desktop and laptop hardware probably don't need to answer 20 questions about privacy controls; so, you might implement a modified version of the schedule, as shown in Table 10.2.

In this scenario, the SIG level 1 would be used only for the moderate-sensitivity vendors, and the low-sensitivity vendors would undergo any formal SIG assessment beyond the basic security questions that would be asked during

Table 10.1 SIG-Based Vendor Schedule—Example 1

Service Sensitivity	First Year	Second Year	Third Year	Fourth Year
High	SIG detailed	SIG level 2	SIG detailed	SIG level 2
Moderate	SIG level 2	N/A	SIG level 1	N/A
Low	SIG level 1	N/A	N/A	SIG level 1

Table 10.2 SIG-Based Vendor Schedule—Example 2

Service Sensitivity	First Year	Second Year	Third Year	Fourth Year
High	SIG detailed	SIG level 2	SIG detailed	SIG level 2
Moderate	SIG level 1	N/A	SIG level 1	N/A
Low	N/A	N/A	N/A	N/A

the vendor selection process specific to the service. However you slice it, the point is that use of a measure like the SIG allows you a lot of flexibility.

Now, imagine a world where you have three versions of the SIG prepared and ready to distribute to your clients and partners immediately upon request. This will never completely replace requests for individual assessments of some services or requirements up front to perform architectural reviews, and so on, but it can reduce the load on your teams significantly. Version 6 of the SIG was released in 2010, and at the time this book was published, it was still unclear how the improvements will affect the adoption rate in the industry. Of course, some clients still may not accept a standardized response like the SIG, no matter how detailed it is, or even a third-party certification; so, you will need to leverage internal risk assessment activities as sources of information when responding to these client queries.

Basing Assessments on Sensitivity

In Chapter 4, risk profiling and risk sensitivity were discussed in detail. This discussion touched on vendor profiling, but didn't get into specific questions that should be included in a third-party profile versus a security risk profile for an internal resource. Concerns around third-party providers are going to focus on a few areas, such as the following:

- Will the vendor store or process sensitive data at their site?
- Will the vendor have access to regulated information?
- Will the vendor's systems directly connect to your organization?
- Will the vendor's service or product be integrated into your offerings?
- Will the vendor's staff need access to your facilities?

These types of questions would be included in a vendor risk profile and then used to determine the sensitivity of the service being provided. You could even include more specific questions about the types of data involved in the service, such as the following:

- Will the vendor store or process sensitive employee data at the vendor location?
- Will the vendor store or process sensitive customer data at the vendor location?

With the introduction of the SIG, version 6, in 2010, there were several improvements made that eliminated one level of assessment detail and, at the same time, made it easier to structure the assessment detail level around the answers to your profiling questions. Version 6 of the SIG includes a SIG-Lite, which is similar to the previous SIG level 1, but no longer includes a SIG level 2 questionnaire for those moderate-sensitivity vendors. Instead, you can use the SIG-Lite as the base for all the assessments and then add individual topic-based questionnaires as needed. For example, you may structure your assessments as follows:

- If the vendor's risk sensitivity rating is High, then the following questionnaires need to be completed:
  - SIG-Lite
  - SIG-F. Physical and Environmental Security
    - required if the vendor will process or store any employee privacy data or client confidential data
  - SIG-G. Communications and Operations Management
    - required if the vendor's system or network will be directly integrated with your environment
  - SIG-I. Information Systems Acquisition Development and Maintenance
    - required if the vendor will have direct or indirect access to your production systems
  - SIG-J. Incident Event and Communications Management
    - required if the vendor will process or store any sensitive data
  - SIG-K. Business Continuity and Disaster Recovery
    - required if the vendor's service will be integrated into your offerings to customers or supports a critical service
  - SIG-L. Compliance
    - required if the vendor will process or store any regulated data at their site
  - SIG-P. Privacy
    - required if the vendor will process or store privacy data for your employees or your clients at their site
- If the vendor's risk sensitivity rating is Moderate, then the SIG-Lite questionnaire needs to be completed.
- If the vendor's risk sensitivity rating is Low, then no further assessment is required beyond the questions in the *vendor security risk profile*.

If you design your security risk profile for each vendor to capture this information from the business owner of the relationship, then you can easily determine the proper level of due diligence required. Of course, you could also create your own questions, but then you are contributing to the problem for service providers who have to respond to so many customized questionnaires. It is better to start with a standard question set and pick and choose which items to include in your subset.

Improving the Process

Having a single vendor assessment format is certainly not the silver bullet; there are several other efficiency improvements that you can make rather easily. The first is to develop a public-facing document that summarizes your security program at a high level, almost like a marketing brochure. This can be a very helpful tool for your sales team to able to provide to new prospective clients/partners before they get into the detailed analysis stage. Include brief summaries of aspects of your program like the general philosophy, alignment with any industry standards, and your high-level privacy policy. Again, this will not replace a detailed analysis later, but it may help to satisfy any concerns that the client's security team will have and help to build confidence that your organization takes security seriously.

Regardless of whether you choose to implement a standardized questionnaire like the SIG or not, you will need some repository for past client questionnaires and answers. If your organization is most often the vendor being reviewed, then you need a way to quickly search old questionnaires for already approved answers or, in the case of repeat client questionnaires, to reuse and refresh the answers from last time. Be cautious about reusing any answers without first viewing them because things change quickly in everyone's environment, and those answers could easily be out of date. Some sort of searchable database of past answers would be a useful investment so that your staff isn't forced to spend a lot of time writing new responses to the same questions. If a database is used as the client response repository, then it needs to be searchable by at least client name, date, and question keywords.

Another option is to align your internal risk assessment processes, such as a Certification and Accreditation review or internal audit function, to the standard assessment questionnaire to ensure that the information is being captured and kept fresh all year round. Otherwise, the task of refreshing it yearly can be significant when you have to break it up into pieces and get each operational team to review and approve their answers. Whenever you can combine assessments, or at least streamline them, you will earn appreciation from the business. If you find yourself in the situation where clients are commonly asking you questions that you have never asked internally, then you will probably want to update your own internal risk assessment process to incorporate these areas of concern. A good indicator of how the threat landscape is shifting is the trends in changing client due diligence focus areas. One year it might be high availability, and the next year, end-point data leakage protection controls. Being aware of these changes in focus can be invaluable when planning and prioritizing your own assessment focus areas.

One final piece of the overall process optimization is feeding any data gathered back into the policy and standard governance process. If your policies or standards seem to be out of line with what clients or partners are expecting, then you should flag these areas and entertain adjusting the internal policy to match.

Similarly, if all your providers are coming up short in your assessments in a certain area, then you might consider adjusting the expectations in your internal policy to allow for some kind of compensating control that provides equivalent protection.

SUMMARY

Whether to use a questionnaire or interview style of assessment can be an important decision that will affect how quickly you get answers back from the other business units and how detailed the responses are. Interviews will provide you the richest information, but the questionnaire is more scalable and less intrusive. Similarly, you should carefully consider at the beginning of every assessment whether you want to use an active testing technique that will produce the most reliable results, or use a less-intrusive passive testing method. When faced with a one-time assessment of a new project, strongly consider Peltier's FRAAP methodology for streamlining the process. Assessments of your program from clients and assessments of your own service providers can be a large resource drain if you don't implement a standardized approach. Try to maximize your resources by producing customer-facing documentation about your security program to minimize ad hoc requests.

Action Plan

After reading this chapter, you should consider taking the following actions in your own organization:

- Pick a few of your most sensitive business areas and schedule an in-person interview with the SMEs to perform the risk assessment, instead of sending them a questionnaire.
- Perform a targeted Google search for any information about your organization that is publically available.
- If you have any active monitoring devices in your environment, including intrusion detection systems or even logs from a firewall, pick a random sampling of data for, say, 30 minutes and review it for any anomalies.
- If you find that you are spending a lot of time assessing and remediating risks associated with the nonsecure configuration systems or software, initiate a project to establish security baselines that meet your standards and focus on automating compliance checks.
- For the next project-based assessment, try the FRAAP approach and worksheets.
- Download the BITS SIG and consider standardizing on it for assessments of your third-party providers and/or make a version available to your clients as a substitute for ad hoc reviews.
- Review your schedule for third-party assessments and ensure that the assessment frequency is directly tied to the sensitivity of the vendor service.

References

[1] Top 100 Network Security Tools. SecTools.Org. http://sectools.org (accessed 20.01.11).
[2] NIST 800-37, Guide for Applying the Risk Management Framework to Federal Information Systems. http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf (accessed 29.12.09).
[3] T.R. Peltier, Information Security Risk Analysis, second ed., Auerbach Publications, Boca Raton, FL, 2005.
[4] SANS Institute, 20 Critical Security Controls for Effective Cyber Defense. www.sans.org/critical-security-controls (accessed 19.05.10).
[5] BITS Standardized Information Gathering (SIG) questionnaire. www.sharedassessments.org (accessed 02.02.11).