

Chapter 3
Conducting a Risk Analysis

Randy Marchany

**Computer and Network Security
in Higher Education**

Mark Luker and Rodney Petersen, Editors

A Publication of EDUCAUSE

Copyright 2003 Jossey-Bass Inc.

Published by Jossey-Bass, A Wiley Company. Reprinted by permission of John Wiley & Sons, Inc. For personal use only. Not for distribution.

Conducting a Risk Analysis

Randy Marchany

A common response to computer security professionals' efforts to secure enterprise and desktop systems is "I don't have anything on my computer that a hacker would want." This statement is true most of the time. Most hackers really don't want your data; rather, they want your computer. They want to use your system to attack other sites. The distributed denial-of-service (DDOS) attacks of the past years are an example of this strategy. Hackers often compromise systems not to steal the data that reside on them but to use the systems to attack other systems. An early DDOS attack involved close to 300 systems used to attack a single site. It's interesting to note that the hackers had complete access to any file on these systems and they could have modified or deleted any files, yet they didn't. How do the hackers get in despite efforts to keep them out?

In 1998, the Internet Audit Project (Siri, 1998) scanned more than 36 million Internet hosts for common and publicized security vulnerabilities. The scan exposed large numbers of vulnerabilities (more than 700,000) to the most common security threats. "These open points of penetration immediately threaten the security of their affiliated networks, putting many millions of systems in commercial, academic, government and military organizations at a high compromise risk" (Siri, 1998, unnumbered). Every vulnerability discovered by the scan could have been eliminated by

proper application of *patches*, or updates to the software that are typically supplied by the vendor. This little test, which took only twenty-one days to run and tested for only eighteen vulnerabilities, showed just how easy it would be to compromise systems in critical industries and the devastating effect that these compromises could have. Ensuring that a university's assets are not vulnerable becomes the primary role of the institution's security officer, system administrators, and internal audit group. Cooperation among the three groups is essential.

Risk Analysis

A critical activity for the security officer, auditor, and IT department to undertake is an *institutional risk analysis*. Conducting a risk analysis is a process of identifying assets, the risks to those assets, and procedures to mitigate the risks to the assets. Individual users and their institutions need to understand what risks exist in their information asset environment and how those risks can be reduced or even eliminated. Embarking on a process to complete such an analysis or self-assessment is critical in today's advanced technological world. The process is one that will benefit both the individual department and the institution as a whole. When conducted in partnership by IT and institutional auditors, a risk analysis can not only provide valuable information to the university but also will carry the weight of a united perspective when it comes time to make decisions about acceptable levels of risk and to secure funding and implement the plan for mitigating risks.

Many good models are available that detail how to perform a risk management study, including how to classify information technology assets and risks. These models provide a foundation for doing more than just avoiding risk—they all provide recommended approaches for identifying weaknesses in the systems, processes for making decisions about how to protect assets, and ways to help evaluate and answer the question, How much security is enough?

Selected Models

The National Infrastructure Protection Center (NIPC) released a document entitled “Risk Management: An Essential Guide to Protecting Critical Assets” in November 2002. This document discusses a five-step risk assessment model: (1) asset assessment, (2) threat assessment, (3) vulnerability assessment, (4) risk assessment, and (5) identification of countermeasure options. The guide includes some examples in tabular format of each of the assessment phases. These examples are useful in helping an organization perform a risk management study. The model is discussed in detail in “The Risk Assessment: Five Steps to Better Risk Management Decisions” (Jopeck, 1997).

Mohammad H. Qayoumi (2002) wrote an excellent guide on continuity planning titled *Mission Continuity Planning: Strategically Assessing and Planning for Threats to Operations*. This booklet was published by the National Association of College and Business Officers (NACUBO) and contains excellent information on risk management, disaster preparedness, business continuity planning, calculating system reliability, and addressing facilities-related risks.

NACUBO also released a document called “Developing a Strategy to Manage Enterprisewide Risk in Higher Education” (Cassidy and others, 2001). This publication presents a definition of risk, the drivers of risk, advice on implementing a risk management plan, and how to advance the risk management agenda to management.

The National Institute of Science and Technology (NIST) published *Security Self-Assessment Guide for Information Technology Systems* (Swanson, 2001). The guide contains a questionnaire template that can be adapted to your site’s needs. The guide is most useful for federal government agencies, since it uses NIST’s federal IT security assessment framework, which standardizes five levels of security criteria.

Case Study: The Star Project

In the late 1990s, a newspaper article about a cyberattack on a university in Virginia prompted a member of the Virginia Tech Board of

Visitors to ask whether a similar attack could happen at Virginia Tech. How vulnerable were Virginia Tech's networks and systems to outside intrusion? This question led to a directive to form a committee to investigate and report on the status of the IT organization's assets. This became known as the STAR (Security Targeting and Analysis of Risks) process (security.vt.edu/playitsafe/index.phtml).

The Information Security Committee was made up of department managers and system and network administrators. This blend of management and technical people was critical in balancing the contrasting viewpoints of security versus access. The prejudices and perspectives of the members resulted in a healthy exchange of viewpoints.

The committee was charged with identifying and prioritizing information systems assets, associating risks with those assets, and listing controls that could be applied to mitigate the risks. Although the committee did not initially consider assets outside of the IT organization, in later iterations of the process departmental assets were evaluated.

Identifying Assets

The committee first compiled a list of division assets and categorized them as critical, essential, or normal assets. An asset was deemed *critical* if the loss of its function would result in the university ceasing to function as a business entity. An *essential* asset would cripple the university's capacity to function, but it could survive for a week or so without the asset. All effort would be made to restore the function within a week. An asset was deemed *normal* if its loss resulted in some inconvenience. The interesting result of this process was that no one wanted to classify his or her asset as normal. The committee avoided trying to classify which components of the network were critical, such as the routers, cable plant, hubs, and switches, by treating the entire network as a critical asset. (These individual network components would have been considered if this had been an audit of the network group.) A sample asset list and classifications are shown in Table 3.1.

Table 3.1. Sample Classification System for Assets and Their Priority.

| Description of Asset | Machine Name | Priority ^a |
|------------------------------------------|-----------------|-----------------------|
| Authentication-authorization services | host1.dept.edu | C |
| DNS name server | host2.dept.edu | C |
| Physical plant, environmental servers | host3.dept.edu | C |
| DNS name server (secondary) | host4.dept.edu | C |
| Network (routers, servers, modems, etc.) | host5.dept.edu | C |
| HR database server | host6.dept.edu | E |
| Payroll server | host7.dept.edu | E |
| Production control servers | host8.dept.edu | N |
| Client systems (Win95/NT, Macs) | host9.dept.edu | N |
| Database group “crash-and-burn” system | host10.dept.edu | N |

^aC, critical element; E, essential; N, normal.

Once a list of assets had been determined and categorized, the committee prioritized them by criticality to the division’s operation. This was done by committee vote. Table 3.2 illustrates a weight matrix that was used to record votes for the assets. The committee votes are recorded in the individual cells, and the total votes for the asset are recorded in the bottom row. The example in Table 3.2 shows that the top three critical assets are the network, the physical plant and environmental servers, and the primary DNS server. Assets were prioritized by voting whether asset 1 was more critical than asset 2, and so forth.

Determining Risk

The committee followed a similar procedure for listing and categorizing the risks to the assets. Four criteria were used in determining a critical risk: (1) It would be extremely expensive to fix, (2) it would result in the loss of a critical service, (3) it would result in heavy, negative publicity, especially outside the organization, and (4) it had a high probability of occurring. Table 3.3 shows the critical risks as determined by the committee. A voting procedure similar to that used to prioritize the assets was used by the committee to prioritize

Table 3.2. Sample Asset Weight Matrix to Prioritize IT Assets.

| | A/A | DNS(p) | Plant | DNS(s) | Network | HR |
|------------------------------------------|-----|--------|-------|--------|---------|------|
| Authentication-authorization services | | 9 | 9 | 4.5 | 9 | 5 |
| DNS name server (primary) | 0 | | 9 | 0 | 9 | 5 |
| Physical plant, environmental servers | 0 | 0 | | 2 | 9 | 4.5 |
| DNS name server (secondary) | 3.5 | 9 | 7 | | 9 | 5 |
| Network (routers, servers, modems, etc.) | 0 | 0 | 0 | 0 | | 0 |
| HR database server | 4 | 4 | 3.5 | 4 | 9 | |
| TOTAL VOTES | 7.5 | 22 | 28.5 | 10.5 | 45 | 19.5 |

the risks. The end result of this entire process was an asset matrix listing every IT asset, an asset weight matrix rank ordering the criticality of the asset, a risk matrix listing every risk, and a risk weight matrix for the set of risks associated with the asset.

Once a final list of assets and risks was developed, the team members mapped the rank-ordered assets and risks into a single matrix. This risk-asset matrix provided guidance as to the order in which each asset and risk were to be examined.

Finally, the STAR team created a controls matrix that listed all of the possible controls that would mitigate the risks listed in the risk matrix. The team did not prioritize these controls; instead it created another risk-asset-controls matrix that listed the possible controls for a particular risk to a particular asset.

Table 3.3. Sample Risk Classification Listing Critical Risks Only.

| Risk | Description |
|---------------------------------|--------------------------------------------------------------|
| Clear text | Clear text data moving among our systems and networks |
| Client system access control | Control of access to distributed desktop client workstations |
| Construction mistakes | Service interruptions during construction, renovations |
| Key person dependency | Too few staff to cover critical responsibilities |
| Natural disaster | Flood, earthquake, fire, etc. |
| Passwords | Selection, security, number of passwords, etc. |
| Physical security (IS internal) | IS private space (machine room, wire closets, offices, etc.) |
| Physical security (IS external) | IS public space (laboratories, classrooms, library, etc.) |
| Spoofing | E-mail and IP address forgery or circumvention |
| Data disclosure | Inappropriate acquisition or release of university data |
| System administration practices | Adequacy of knowledge, skills, and procedures |
| Operational policies | Appropriate strategies, directions, and policies |

Applying Controls

Risk and asset matrices formed a blueprint for applying controls to the assets. The STAR team developed a set of compliance matrices that corresponds with the risks and assets listed earlier. This set of compliance matrices contains detailed line-item actions to verify that a particular task has been performed. A color-coding system is used to denote success or failure of a particular line item. This provides an auditor, security manager, or system administrator with a quick way to verify the compliance of an asset with the risk analysis process.

Figure 3.1 shows a portion of the original overall compliance matrix, which is similar to the executive summary portion of an audit. It lists an overall rating for each critical asset. The risks are listed along the y-axis and the assets are listed along the x-axis. There is a corresponding set of matrices associated with each risk line item that contains more detailed information on the tests required to determine the vulnerability of an asset to the particular

Figure 3.1. Summary Compliance Matrix Showing the Overall State of an Asset in Relation to Identified Risks.

| | | IS ASSETS | | | |
|---------------------|-------------------------|-------------|-------------|---------|--------|
| | | Site 1 | Site 2 | Site 3 | Site 4 |
| UNIX security risks | OVERALL | OK | | | |
| | System admin. practices | OK | OK | | |
| | Data disclosure | FAIL | OK | CAUTION | |
| | Passwords | OK | CAUTION | OK | |
| | Key person dependency | OK | FAIL | FUTURE | FUTURE |
| | Physical security | CAUTION | CAUTION | OK | FUTURE |

risk. A set of detailed commands needed to check the system is at the lowest level. The test procedure is as follows:

1. Use the detailed command list to perform the test.
2. Record the results in the detailed compliance matrix.
3. Compute an overall score for an asset (70 percent “OK” means an “OK” score at this level, for example) and record the score in the overall compliance matrix.

The matrices allow Virginia Tech staff to track progress in addressing the risks over the long term. They provide a foundation on which to apply the scarce resources to perform a cost-benefit analysis of the assets and risks.

After conducting analyses with the IT organization, the STAR team took the matrices and the process and made them available to every department in the university. STAR provided the common risks set to departments and encouraged each to add its own risks. The end result is that each department can now look at the risks it has in common with others and at its own unique risks about which it needs to be concerned.

STAR Today

The STAR methodology is still in use at Virginia Tech today, and it continues to be refined, even after seven years. It took the committee about one year of meeting once every two weeks to develop and refine the methodology. Subsequent risk analysis projects are much less time consuming, since the basic matrices now exist. In 2002 the STAR team completed the seventh iteration of the risk analysis process in one month. Support from university leadership has also been critical. Securing the approval of the top management at Virginia Tech ensured a 98 percent on-time return rate of the individual departmental risk analysis reports. This simply confirms

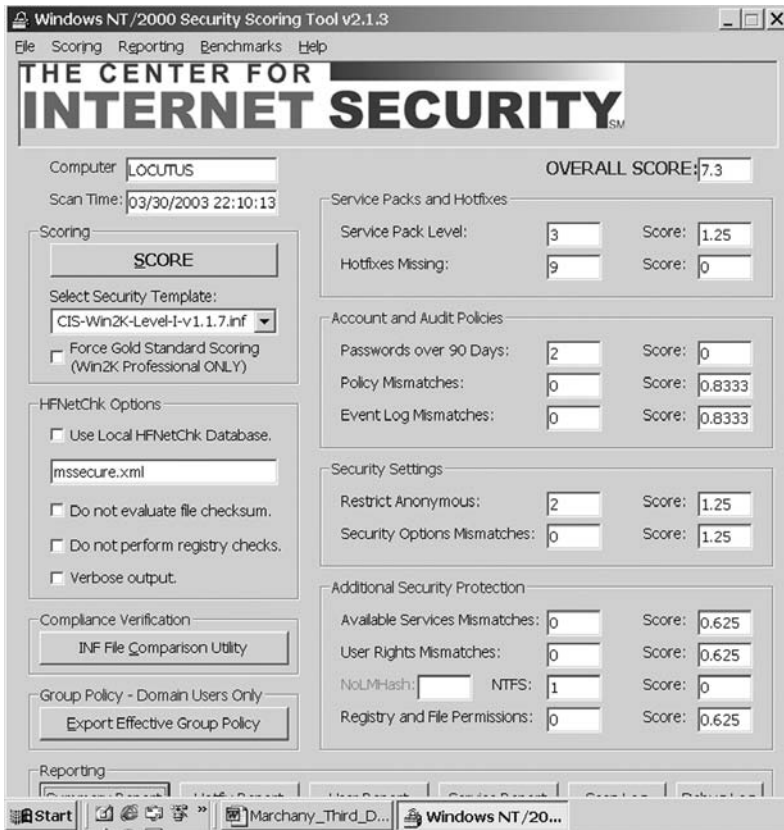
what every security office has known: management buy-in of the process is critical to its success.

Although STAR's main purpose is to provide a repeatable method for prioritizing assets and risks to those assets, one of the major modifications to the process during the last few years revolves around asset classification. Initially, computer systems or the network were identified as assets, but the team is now moving toward identifying business processes as assets. This, in turn, creates a layered approach to asset classification. For example, at the top level, student registration is identified as a business process and asset. It comprises (1) departments that manage the business process, that is, the registrar; (2) the software management group that manages the software that runs the student registration process, for example, SCT's Banner software; (3) the information systems group that manages the machines that run the Banner software; and finally, (4) the actual machines that run the Banner software. The STAR methodology can then be used to prioritize assets at any level of this tree. Virginia Tech is now using this layered approach to help identify key business processes, thereby in turn helping design a better disaster recovery procedure. An additional benefit to "layering" the assets is being able to see the dependencies that a particular asset needs to accomplish its mission.

The checklists used to measure compliance have also changed, and Virginia Tech has adapted the STAR methodology to use the Center for Internet Security's (CIS) (www.cisecurity.com) security benchmarks (replacing the colored matrices) to measure how a computer system complies with local security requirements. The CIS benchmarks are free and prove to be an excellent resource for various UNIX, Windows, and router platforms. The CIS benchmarks provide a straightforward method of configuring a system according to the STAR analysis. CIS benchmarks are applied to the critical assets. The CIS toolkit provides a scanning tool that rates the asset based on the benchmark. This score presents a target value for auditors.

Figure 3.2 shows the output screen of the CIS Windows 2000 scanning tool. When the necessary services are enabled on a critical asset, the scanning tool reports a score for each service. This scoring method provides an auditor with a simple way of checking an asset for compliance. If the asset scores a 7.3 or higher, then its local policies have been set in compliance with the risk analysis. A lower score requires justification from the asset system administrators or owners. It does not mean the system is less secure; instead, it generally means that something required the system security policies to

Figure 3.2. Sample CIS Scanning Tool Output.



be relaxed. In most cases, vendor software requirements are the primary reason that system security policies are relaxed.

Conclusion

Performing a risk analysis is a necessary first step in assuring the security of campus technology resources. A partnership between IT and auditors on campus acknowledges the common goal of these two groups and can be effective in garnering support for implementing plans to mitigate risks.

Many resources are available to help institutions perform a risk management study. Institutions should select a method to identify and categorize assets and the risks to those assets, and use a simple, replicable process, such as the STAR process, to prioritize assets and risks.

The matrices and scoring tools described or referenced in this chapter provide a quick way for auditors to determine compliance with the institution's standards. The same matrices give systems administrators the status of security controls installed on a particular asset. As an example, the STAR technique simplifies standard risk analysis methods and makes it easier for all departments to provide meaningful information to the institution.

As with other risk management activities, conducting a security risk analysis is not a one-time event. Audits should be performed regularly to ensure compliance with critical security measures and with plans for mitigating risk. Periodic evaluations of the entire IT security program are necessary to maintain agreed-on levels of security.

References

- Cassidy, D., and others. "Developing a Strategy to Manage Enterprisewide Risk in Higher Education." Washington, D.C.: National Association of College and Business Officers, 2001.
- Jopeck, E. "The Risk Assessment: Five Steps to Better Risk Management Decisions." *Security Awareness Bulletin*, 1997, 3-97, 5-15.

- National Infrastructure Protection Center. "Risk Management: An Essential Guide to Protecting Critical Assets." [www.in.gov/c-tasc/whatsnew/risk_management11-02.pdf]. Nov. 2002.
- Qayoumi, M. H. *Mission Continuity Planning: Strategically Assessing and Planning for Threats to Operations*. Washington, D.C.: National Association of College and University Business Officers, 2002.
- Siri, L. "Internet Audit Project." [www.viacorp.com/auditing.html]. 1998.
- Swanson, M. *Security Self-Assessment Guide for Information Technology Systems*, NIST Special Publication 800-26. Washington, DC. National Institute of Standards and Technology. [csrc.nist.gov/publications/nistpubs/800-26/sp800-26.pdf]. Nov. 2001.