The "Zen" of Risk Assessment

The time and resources needed for proper risk assessment can be mitigated and the benefits magnified by making assessment an ongoing, rule-based process

By **Cedric Bennett** and **Richard Jacik**

Now, here *you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!*

—Lewis Carroll
Through the Looking Glass[1]

Whenever a person decides to act (or not) to repair a computer security vulnerability, prevent an Internet-based attack, or introduce a procedure designed to support an information security initiative, assessment of risk has taken place. Perhaps the individual was not specifically aware of performing a risk assessment—it might have been mostly implicit and possibly of very narrow scope—but some assessment did nevertheless take place. At the very least, some amount of thought went into the importance of the risk being mitigated and the potential for success of the mitigation chosen. As a practical matter, this sort of information-security decision making goes on every day, perhaps multiple times, for every person responsible for any aspect of information security.

Although this simple case illustrates an example of risk assessment, the term is usually reserved for the assessment of risk from some larger, organizational perspective. As pointed out by Stoneburner, Goguen, and Feringa, "The principal goal of an organization's risk management process should be to protect the organization and its ability to perform [its] mission...."[2] Put another way, wrote Alberts and Dorofee, a risk assessment should "look at the organization itself and identify what needs to be protected...."[3]

Risk assessment and risk management cover a landscape far larger than information technology. The entire insurance industry is built around those concepts and obviously is concerned with a broad array of subjects (such as property, plant assets, and equipment, as well as individual life and health). This Good Ideas article focuses specifically on the cybersecurity risks associated with information technology, systems, and services.

Traditional Risk Assessment

An organization that wants to protect its information assets appropriately will perform a risk assessment. Since it is axiomatic within information security circles that complete security is unattainable (or unaffordable), risk assessment becomes, in part, a process of deciding which assets should receive the highest levels of protection and which can tolerate lower levels of protection. Because risks, vulnerabilities, processes, and technology all change over time, risk assessment is also usually thought of as a cyclic process.

Experts agree that risk assessment is a crucial early step in effectively developing and implementing an information security program. Risk assessment aims to assist development of a security strategy and the plans to carry out that strategy with the goal of mitigating the risk to the organization's most critical resources. Traditional approaches to risk assessment employ a very rigorous and comprehensive process. Unfortunately, they usually take considerable time to complete and sometimes produce plans that themselves seem overwhelming. Given the large investment in time and resources, a traditional risk assessment can also beg the question of what to do about risk mitigation activities while the long-running assessment project is in process.

Most information security practitioners agree that traditional risk assessment is a very large task (many find it quite daunting). The *EDUCAUSE Effective Security Practices Guide*[4] suggests breaking this project/process into three phases: Preliminary Risk Assessment, Risk Analysis of Critical Areas and Processes, and Institution-Wide Risk Assessment. The first phase alone is estimated to take from four to six months, according to the guide, to achieve the objective of "... giving upper management a concrete overview of the IT risks leading to more resources being allocated to address major problems." Is it any wonder that according to the 2003 EDUCAUSE Center for Applied Research (ECAR) survey on security, only 30 percent of higher education institutions surveyed had conducted a risk assessment?[5]

Leverage What You Know

We recommend an approach for assessing risk in which overall risk assessment is more of an ongoing process than a project. It produces usable

results from the start, which can provide broad guidance for security strategies and plans and also focus traditional risk assessment toward specific assets and resources.

Develop Data Classification Rules[6]

The process starts by first developing a brief set of basic rules for determining the criticality of different types of data from a risk perspective. One effective approach establishes rules that divide all data into three levels of general risk criticality—most critical, critical, and least critical. This one-time effort should produce only a few necessary rules that are broad in nature and easy to understand and apply.

Developing such rules is easier than it might seem because general data types in higher education are similar from institution to institution and because every institution already understands most of the rules—usually they just need explicit documentation. For example, one rule would probably focus on the institution's legal data requirements. Such a rule might declare that any data protected by federal, state, or local regulation fits into the most critical category and that any data protected by contractual commitment fits into the critical category. Another rule might specify that all data providing information regarding access to resources (for example, password files, access authority files, building-key information, and physical-plant data) falls into the most critical category (except for campus maps and other such information provided for public access, which would fall into the least critical category). Other likely rules would focus on intellectual property, financial data, and data about individuals.[7]

The rule set developed should be relatively complete in its first implementation, but it need not be exhaustive. Later steps in the process will make any missing data types obvious; at that time any necessary additional rules can be added to the set. Because these rules are based upon already existing knowledge, the initial set can be developed very rapidly, usually in one to two weeks.[8]

Apply the Rules to Classify Data Collections and Related Resources

Data are usually kept in collections called databases, files, tables, and others. In most data collections, more-sensitive data elements are rarely segregated from less-sensitive ones. When determining the data classification level of any collection, the classification of the most critical data in that collection determines its classification. This is actually helpful, since it means that classification occurs primarily by inspecting the collection at an aggregate level. For instance, if a data collection includes patient health information, it falls under the federal regulations described in HIPAA. Therefore, the entire data collection, including related data that is not patient health information, falls into the most critical classification. Under these circumstances, applying the rule set is fairly straightforward.

You can look at related information resources following the same aggregation approach. Assign the most-critical category to any information system that processes most-critical data collections. In the same way, servers, network segments, and any other information resources that support the most-critical system and its data are themselves categorized as most critical. Follow the same pattern for the critical and least-critical categories.

Unlike with traditional risk assessment approaches, you do not need to inventory all assets before classifying the various levels of criticality. With this rule-based approach, all or most of the critical assets become obvious immediately. If it happens that a critical asset is overlooked initially, application of the rule set makes its classification obvious when it is eventually uncovered. Moreover, if the objective is to provide upper management with information, this process can get an institution to that point in much less than four to six months.

Once the classification process is complete, the institution's assessment of risk has taken a large step forward. If the formal process goes no further, at least the criticality of information resources has been identified. That information can be used in any subsequent decision-making process with regard to effort, expenditure, and focus for information security. We recommend, however, that institutions following this approach take at least one additional step.

Develop Broad Strategies

It is not necessary to exhaustively inventory and classify all data collections before proceeding, although it is best if most data types have been identified and classification rules developed. Then, as any new collections are uncovered, existing rules can be quickly applied, providing immediate guidance.

Moreover, rather than focusing initially on specific risk issues of particular information resources (and because this process deals with broad categories of information assets), a more effective approach might be to start at a strategic level. If an institution is just beginning to develop an information security program or hasn't yet applied many risk mitigation procedures or technologies, starting from a more strategic perspective will likely provide a greater degree of overall protection much faster. Even if an institution has already deployed some information security resources, the strategic approach will help ensure that major areas have been covered effectively.

Once the institution has identified multiple information assets as most critical, it can begin to identify threats and vulnerabilities that generally exist across that set of resources. "Threats" can be thought of as something that might negatively affect an information asset, represented by the loss of confidentiality, integrity, or availability of that asset. "Vulnerability" can be thought of as some weakness that will allow the asset to be exploited—the "how" of the threat.

The object of this step is to deploy risk mitigation processes and technologies that apply across a large set of assets, whether focused on most critical assets or all information assets. The important benefit is that at least some degree of protection is put in place and leveraged across a large set of resources. Thinking about threats and vulnerabilities

at this level is best accomplished by thinking both in terms of broad threat classes such as physical, network level, host computers, and applications on the one hand and sources of threats such as inside and outside people and hardware and software defects on the other hand.[9] Once this sort of list exists, general risk mitigation strategies can be applied.

It should be obvious, for example, that any resources accessible to the Internet risk malicious outside attack on a regular basis. The application of firewalls in front of information resources is an obvious mitigation strategy. Applying the notion that more critical resources ought to receive the most risk mitigation, a strategy might be developed to apply multiple layers of firewall around the most critical and critical resources wherever possible. Similarly, since any desktop or laptop computer can be at even greater risk of Internet-based attacks, a strategy might be developed that requires any such computer containing most critical and critical data to apply additional protection. A rigorous patch-management system might be required, for example, or perhaps all critical data resources must be encrypted on these computers.

Once these risk categories are assigned for a major subset of data, systems, and associated resources, strategy development and security planning can begin. Strategies can be developed for the different levels of criticality, and plans can allocate information security resources in appropriately proportionate ways—most to protect most critical and critical assets, and less to protect least critical assets.[10]

This level of risk assessment can also lead to a more constructive executive and budgetary dialogue on information security resource requirements. As plans are developed, vulnerability and threat assessment can proceed following whatever approach seems most appropriate for the institution and the issues addressed.

Once these general strategies are turned into policy and begin to be applied, the institution will know that it has begun to protect itself against the most likely vulnerabilities and risks. At this point more traditional risk assessment approaches can be applied effectively, first focusing on the major collections in the most critical category. This narrowing of focus ensures that the institution's limited resources are applied in the most cost-effective way and that, at a minimum, processes are in place to protect the most important resources.

Additional Benefits

In addition to the major advantages already described, this approach provides some other, less obvious benefits.

System Development

Once in place, the risk assessment rules can be applied to developing systems and proposals for new systems and technologies. New implementations that fall into the most critical risk category can be more carefully tracked and audited from the beginning to ensure compliance with regulations and institutional policies.

System Development Life Cycle

These ideas can also be built into the institution's system development life cycle. Systems at the higher end of the risk spectrum might have more reviews scheduled during development and implementation and more frequent audits after implementation.

Extension to Other Areas of Risk

This rule-based approach to cybersecurity risk assessment is an easy concept to teach and build upon. As new considerations become important, it becomes an extremely useful tool for determining required levels of protection. For example, are there any issues with regard to the use of personal computers as distributed nodes in the deployment of a new information system? The answer is informed by asking questions about the risk category of the data that will end up residing on those computers (however temporarily). If the data belongs in a high risk category, then additional mitigation must be applied to the personal computers participating in the deployment.

Information Security Awareness

The rules can be written easily and shared widely across the institution. They can even be reduced to a simple table, with examples. As information system and asset inventories take place (whether for risk assessment or other purposes), the rules can quickly be applied to provide up-to-date risk assessment information. Not only does this work inform the information security strategy, it also helps in the ongoing security awareness process.

One major university that implemented a version of this approach simply created three categories, A, B, and C, for most critical to least critical, respectively. After some time had elapsed during which these concepts were used repeatedly, the terminology entered the common vocabulary of both information system professionals and their clients. Before long, project proposal discussions were regularly laced with references like "Category A data," "Category C desktop computers," and "Category B servers"—an unanticipated but welcome additional outcome.

Conclusion

A practical approach to dealing with cyber risk entails finding sensible ways to achieve necessary improvements in the short term, even if that means briefly delaying a fully substantive and exhaustive risk assessment. This article describes one approach that has succeeded on multiple occasions. Appropriate trade-offs between a perfect risk assessment that might never be complete (or even begun) and a very good one that can be implemented quickly and with useful results is a discussion likely to benefit any campus. *e*

Endnotes

1. L. Carroll (Charles L. Dodgson), *Through the Looking Glass*, Chapter 2, in *Logical Nonsense: The Works of Lewis Carroll*, P. C. Blackburn and L. White, eds. (London: The Macmillan Company, 1934), p. 177. First published in 1872.
2. G. Stoneburner, A. Goguen, and A. Feringa, "Risk Management Guide for Information Technology Systems: Recommendations of the National Institute of Standards and Technology," NIST, Technology Administration, U.S. Department

of Commerce, Special Publication 800-30, U.S. Government Printing Office, Washington, D. C., 2001.

3. C. Alberts and A. Dorofee, *Managing Information Security Risks: The OCTAVE℠ Approach* (Boston, Mass.: Addison-Wesley, 2003), p. xxv.
4. See <http://www.educause.edu/security/guide/>.
5. R. B. Kvavik and J. Voloudakis et al., *Information Technology Security: Governance, Strategy, and Practice in Higher Education* (Boulder, Colo.: EDUCAUSE Center for Applied Research, Research Study, Vol. 5, 2003).
6. Information Methodologies, Inc. (http://www.infometh.com), in 2000 developed the proprietary information security methodology underlying much of the discussion in this article. The methodology is not publicly available.
7. Data can also be classified as the result of the application of "prudent stewardship," where there is no reason to protect the data other than to reduce the possibility of harm or embarrassment to individuals or to the institution.
8. See <http://securecomputing.stanford.edu/dataclass.html> for a description of one institution's rule set and <http://www.stanford.edu/group/security/classification/classification_of_data.html> for the resulting table and lists.
9. An excellent source for questions that can be asked to aid in this step of the process can be found in the "Information Security Governance Assessment Tool for Higher Education" (particularly the section on Technology) published by EDUCAUSE at <http://www.educause.edu/ir/library/pdf/SEC0421.pdf>.
10. Even least critical and public information resources require some measure of protection—at least from accidental or malicious damage.

Cedric Bennett (Ced.Bennett@stanford.edu) is Emeritus Director, Information Security Services, at Stanford University in Stanford, California, and still spends considerable time helping Stanford and other institutions address their information security requirements. Richard Jacik (jacik@infometh.com) is President and cofounder of Information Methodologies, Inc. (IMI), a higher education systems integrator based in Sterling, Virginia, and has worked with educational institutions for 17 years.