Risk Management and the Cybersecurity of the U.S. Government

Input to the Commission on Enhancing National Cybersecurity

Steven B. Lipner and Butler W. Lampson

Executive Summary

Risk management is a fundamental principle of cybersecurity. It is the basis of the NIST Framework for Improving Critical Infrastructure Cybersecurity. Agencies of the U.S. Government certify the operational security of their information systems against the requirements of the FISMA Risk Management Framework (RMF). The alternative to risk management would presumably be a quest for total security – both unaffordable and unachievable.

However, cybersecurity risk management is hard. It is impossible to measure precisely either the amount of security a given investment buys or the expected consequences of less than perfect security. Thus, decision makers must make security investments whose benefits are very uncertain. This makes it tempting to spend less on security and more on new programs or other alternatives with more visible benefits.

The RMF requires an agency to implement security controls specified in NIST SP 800-53 for the sensitivity (low, moderate, or high) of the agency's systems. There are 256 top-level controls and many variations, each described in English text that requires expert interpretation. If a system doesn't implement a particular control, an agency executive must accept the associated risk – but will probably be unable to comprehend either the control requirement or the risk. This vagueness and complexity means that the system is likely to be approved regardless of the actual risk.

Despite the investment in creating the RMF and certifying systems, critical government information systems continue to be subject to successful attacks. We recommend that the government fundamentally reform the RMF in two ways:

- Replace the SP 800-53 controls with a set of fewer than 30 clear requirements that system owners must implement – and are able to understand. Each requirement should have an associated criticality rating (low, medium or high) for cases where an agency or information system cannot meet the requirement.
- Require that when a requirement cannot be met, the decision to accept the risk must be reviewed by both an agency manager who is responsible for accepting the risk, and a peer level manager from a security expert agency who is responsible for ensuring that the agency manager has a clear understanding of the consequences of failing to meet the requirement. The levels of the reviewing managers should be determined by the criticality rating associated with the requirement; the highest criticality risks must be accepted by an agency CIO.

These recommendations are based on our experience with managing the risks associated with secure cloud operations and secure software development at a large commercial vendor, and on our understanding of Australia's successful cybersecurity risk management program. We believe that their implementation would be a significant step in the U.S. Government's progress toward leadership in cybersecurity.

Risk Management and the Cybersecurity of the U.S. Government

Input to the Commission on Enhancing National Cybersecurity

Steven B. Lipner and Butler W. Lampson

Introduction

Cybersecurity is a complex and multi-faceted issue, but this paper focuses on cybersecurity risk management for United States Government systems. We believe a new that approach to risk management is feasible and can lead to significant improvements in the cybersecurity of government systems and information – to the point that the government can become a leader and a positive example of sound cybersecurity practices.

Risk Management

Since the original release of OMB Circular A-130 [1] in 1985, government cybersecurity policy and practice has been based on the principle of risk management: investment in security measures should be no greater than the expected harm prevented (the value of the harm resulting from a successful attack or accident affecting an information system times the likelihood of such a successful attack). In theory, this principle is perfectly sensible: there is no point in devoting more resources to security than the value of the loss resulting from a security failure. Many businesses such as investment management and insurance rely on this principle.

For cybersecurity, however, this doesn't work. A fire insurance company can calculate the expected frequency and consequences of house fires and determine what fire safety measures (such as compliance with building code and inspection of electrical wiring and natural gas supplies) will prove cost-effective, but it is impossible to make such a precise risk determination for cybersecurity:

- We have a poor understanding of adversary capabilities and resources.
- We have a poor understanding of the effectiveness of cybersecurity countermeasures.
- We have a poor understanding of the consequences of a successful attack – in particular, of the extent to which exploitation of a single vulnerability can result in widespread damage to confidentiality, integrity, or availability.

Improving our understanding in these areas is an open research challenge, but there is no reason to believe that it will be feasible to reach the level of precise risk estimation associated with physical domains such as fire or flood. In any case, such precision will not be achieved in the near future.

Government Risk Management Practice

The NIST Cybersecurity Framework [2] incorporates risk management as one of its core principles and expects that adopters of the framework will practice risk management, but leaves most specifics of risk management practice up to the individual organization. Today's U.S. Government cybersecurity risk management practices are based on FISMA (the Federal Information Security Management Act) [3] and on the associated Risk Management Framework (RMF) [4], also developed by NIST.

Under FISMA and the RMF, agencies categorize their systems as low, moderate, or high sensitivity based on an assessment of the sensitivity of the information processed and the potential consequences of loss.

Once the category has been determined, agencies are expected to implement security controls required for that category as specified in NIST Special Publication (SP) 800-53 [5]. SP 800-53 specifies some 256 security controls spanning people (training and management), processes for operation and development, and technology. Many of the controls specify multiple security requirements and some include numerous enhancements that add requirements for special situations. The descriptions of the controls are provided in English text, and thus each is subject to interpretation or misinterpretation in the process of translation into decisions about system design, implementation, configuration, or operation.

Federal information systems are required to comply with the requirements of the RMF, and in many cases agencies assure their compliance by contracting with consultants. These consultants review systems' features and management practices against the list of controls in SP 800-53 and produce lengthy reports documenting compliance. Then an approving official from the agency is expected to review the report and approve the system's operation. Approval may include accepting residual risk resulting from requirements that were not met. One of the authors has discussed agency security with Federal IT security managers and observed that the discussion at least as often turns to the documentation associated with RMF compliance as to threats and (technical or operational) countermeasures.

The ultimate test of a cybersecurity program is how well it protects systems and information, and current government practices are not passing this test. The intrusion into the Office of Personnel Management (OPM) [6] is the most visible recent example of a failure of cybersecurity in the Federal government, but in recent years the press has reported other intrusions or abuses as well – of IRS, of the State Department, of the Pentagon to cite a few. An organization called Security Scorecard has published a comparative analysis [7] of government agencies' and commercial industry security, apparently based on indicators visible from the Internet, and rated government lowest among industry sectors.

The Challenge of Cybersecurity Risk Management

The real-world history of security intrusions into government systems makes it evident that some aspect of U.S. Government cybersecurity risk management is not working well. It is clear that operating large IT systems securely is not easy, and that attackers have an advantage given the numerous ways of attacking systems, but cybersecurity programs must be built to deal with these realities. U.S. Government systems continue to be attacked successfully, and independent observers do not cite government practices as an example of best practices for cybersecurity. Why do problems continue to arise, and is there a better approach to risk management?

Ralph Langner and Perry Pederson [8] believe that cybersecurity risk management is "bound to fail." In summary, their view is that managers have an incentive to underinvest in security measures: while the costs of security measures that mitigate risk are clear, the expected losses associated with failures of security are anything but. If decision-makers insist on committing resources to security, they have less budget available for other programs or activities, and there is no guarantee that the (security) result would have been different if they spent less. While there is no reward for preventing attacks, there are clear rewards for spending resources on new program activities rather than security.

The structure of the RMF exacerbates the problem. The enormous catalog of fine-grained security controls is confusing even to security experts. The size and structure of the catalog drives teams charged with security to focus on convincing themselves that they have met the requirements for individual controls rather than focusing on the overall protection of the system. In addition, whether a control has been implemented is not black and white but rather a matter of interpretation, and it's easy to find interpretations that minimize cost and operational impact, but also minimize security. The RMF process requires Federal executives to approve security programs based on compliance documents and the text of SP 800-53, and to accept security risks that depend on people, processes, and technology that they often do not understand. The inevitable result is a tendency to choose cheaper security measures, regardless of the actual consequences for agency security.

Nonetheless, it is clear that any practical approach to managing cybersecurity for the Federal government – or any other organization – must be based on risk management. The challenge is to ensure that cybersecurity risks – which cannot be quantified accurately – are realistically reflected when organizations make their decisions to implement or omit security measures.

Alternate Approaches to Managing Risk

In their analysis, Langner and Pederson focus on information systems supporting critical infrastructure and argue that policies should specify mandatory security measures: "Fix the design vulnerabilities rather than hypothesize about threats." In effect, they propose doing away with risk management. This has the appeal of simplicity, but it's impractical because it can lead to spending unbounded amounts of time and money on security measures at the expense of system function and usability.

The Australian Signals Directorate (ASD), the information security agency of the Australian government, publishes an information security policy [9] and best practices that in some ways are similar to the RMF. In particular, the Australian government takes a risk management approach. A key difference between the Australian approach and the RMF, however, is that the Australian government has established a very specific set of "Strategies to Mitigate Targeted Cyber Intrusions" [10] that are both reflected in policy as mandatory requirements and supported by specific technical guidance. ASD reports that the "strategies" prevent over 85% of targeted intrusions. Australian government agencies manage risk, but failure to comply with mandatory requirements is much less subjective than is the case with the RMF, and any such failure must be reported to the government agencies responsible for information security.

Commercial software vendors that seek to build products and operate cloud services securely face similar challenges: the effects of secure development and operational practices are not readily measurable, and the market tends to incentivize development team leaders to build customer-visible features rather than improve products' resistance to attack. The authors are familiar with the secure development program at Microsoft (the Security Development Lifecycle or SDL [11]) which has been successful at improving software security. In the SDL, development groups must meet specific technical requirements. Developers receive training on secure development practices, and in most cases, (mandatory) automated tools present developers with very specific indications of security errors that they have made and problems that they must address. With growth in the importance of cloud services, Microsoft created a similar program, Operational Security Assurance or OSA [12], to focus on the operational security aspects of those services. The principles of OSA are similar to those of the SDL: clear requirements, training, and use of automated tools to ensure that security requirements are actually met.

Under the SDL and OSA, if a product team is unable to meet a requirement, a manager from the product team may approve an exception and accept the associated risk, but there are rules to ensure that the risk decision is appropriately informed:

- The level of manager who can approve an exception is determined by the potential harm that could result from accepting the risk. If a decision to accept risk would allow software to ship or a service to operate with a serious vulnerability, an executive-level manager (a corporate vice president) must make that decision.
- Before a manager can approve an exception, he or she must review the exception and the risk it represents with a peer-level manager from the security team that manages the SDL or OSA process and requirements. The security team manager is part of a separate organization, and is responsible for making risk and consequences clear to the product team manager. Because it is not possible to quantify cybersecurity risks, the review is qualitative and draws on the history of product vulnerabilities and attacks that resulted from similar cases to the one under discussion. This sort of informed discussion of risk very frequently results in the product team manager making a decision to deny the exception and insist that the problem be fixed or mitigated.
- Approved exceptions are tracked so that the issue that resulted in the exception can be addressed in a future version of the software or service.

Experience with the SDL and OSA has shown that the processes encourage sound decisions with regard to security without impeding the success of the business. All the authors know about the Australian ASD system is what is on the government websites, but that system appears to support a risk management approach similar to that of the SDL and OSA.

An Alternative for the U.S. Government

If the U.S. Government seeks to improve its cybersecurity practices and become a good example for the nation, we believe that fundamental changes to the RMF will be required. In particular:

- NIST should replace the control catalog of SP 800-53 with a set of mandatory security requirements (MSRs) for government agencies that are clear, concrete, and specified at a level corresponding to actions and decisions that system designers and operators will understand. The ASD controls represent one possible starting point for developing such a set of requirements, and the SANS/Center for Internet Security “Top 20” Critical Security Controls [13] another. MSRs could be tiered by information or system sensitivity, or there could be a single common set. Our suggestion is to create an initial set of baseline MSRs and then augment them as appropriate for more sensitive systems. In any case, each MSR should be characterized in terms of criticality of the impact if it is not met.
- As part of its adoption of a new set of MSRs, NIST should require that software developed by or for the U.S. follow a process similar to the SDL. While SP 800-53 makes passing reference to some elements of a secure development process, government software acquisition contracts (and in-house development efforts) only infrequently implement such a process.
- The government (NIST, DHS, and/or DoD) should obtain a set of automated tools that will enable agencies to deploy and/or verify the implementation of the MSRs. Because the U.S. Government uses only a few software platforms across many agencies, an investment in such

tools would be highly leveraged, unlike the current requirement that each agency or system conduct a manual assessment against the list of controls in SP 800-53.

- OMB should restructure the way that agencies review and approve decisions to accept cybersecurity risks, so that the level of approval for an exception is determined by the criticality of the requirement. One of a few agencies that are experts in cybersecurity should have responsibility for peer review of risk acceptance decisions – NIST and DHS are logical choices in for non-National Security Systems, and NSA for National Security Systems. The reviewing official from the expert agency should be a peer (comparable rank) of the official authorized to accept the risk, and the review should clearly explain the potential consequences of failing to meet the requirement. Because the security requirements under the proposed regime will be specific and associated with specific criticality levels, it should be possible to mandate such peer review without imposing an overwhelming workload on the agencies doing the peer reviews. The combination of clear requirements and peer review should result in better-informed, and ideally less frequent, risk acceptance by agency officials.

The authors believe that adopting risk management practices that are scalable, comprehensible, and have proven successful in other environments will improve the security of U.S. Government information and systems, and thus help the U.S. Government become a good example of cybersecurity practices for the private sector and improve the confidence of citizens in the security of their information.

About the Authors

Steven B. Lipner retired from Microsoft in 2015 as partner director of software security. He is the creator and long-time leader of the Microsoft Security Development Lifecycle (SDL). The SDL was the first scalable and effective approach to achieving security assurance for large-scale software systems and has been applied by Microsoft and other development organizations. Prior to joining Microsoft, he worked in the areas of high assurance secure systems, cryptographic protocols, and Internet firewalls. He was elected in 2015 to the National Cybersecurity Hall of Fame.

Butler W. Lampson is a Technical Fellow at Microsoft. He has worked on computer architecture, local area networks, laser printers, operating systems, programming languages, fault-tolerant computing, computer security, WHSIWYG editors, and tablet computers. He was one of the designers of the Alto personal distributed computing system, the SDSI/SPKI system for network security, and many other systems. He is a member of the National Academy of Sciences and the National Academy of Engineering. He received the National Computer Systems Security Award in 1998, the Turing Award in 1992, and the National Academy of Engineering's Draper Prize in 2004.

References

1. https://www.whitehouse.gov/sites/default/files/omb/assets/OMB/circulars/a130/a130revised.pdf
2. https://www.nist.gov/cyberframework
3. https://www.dhs.gov/fisma
4. http://csrc.nist.gov/groups/SMA/fisma/framework.html
5. http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf
6. https://www.opm.gov/cybersecurity/cybersecurity-incidents/
7. http://info.securityscorecard.com/2016-us-government-cybersecurity-report
8. http://www.langner.com/en/wp-content/uploads/2013/06/Bound-to-fail.pdf
9. https://www.protectivesecurity.gov.au/informationsecurity/Pages/Information-security-core-policy.aspx
10. http://www.asd.gov.au/infosec/top-mitigations/mitigations-2014-table.htm
11. https://www.microsoft.com/en-us/sdl/
12. https://www.microsoft.com/en-us/download/details.aspx?id=40872
13. https://www.cisecurity.org/critical-controls.cfm