The Risk of Propagating Standards

Matt Bishop
Dept. of Computer Science
University of California, Davis
Davis, CA, USA
bishop@ucdavis.edu

Candice Hoke
Cleveland Marshall School of Law
Cleveland State University
Cleveland, OH, USA
shoke@me.com

Information technology is now used pervasively in mission-critical governmental, commercial, and academic processes throughout the world. As with any integration of new technology into existing processes, the goal of disrupting the existing process minimally introduces risks in the use of the new technology, and of the combination of new and existing technology. This is particularly true when the new technology is widely used in other environments that have very different requirements.

"Best practices" documents aid this process. They provide guidance for the initial set-up, configuration, use, management, and decommissioning of these computers. They also guide the process of developing policies and procedures to govern the use of the machines. In the U.S., many best practices are promulgated by the federal government, although system and software vendors may also supply (much more limited) guidance.

Underlying these guidance documents is an implicit assumption that the security methods and techniques that function well for the federal government will also work well at the state and local governments and other entities. Further, transitioning these federal practices into other contexts is seen as routine: simply apply the guidance in the new environment; no significant differences exist that might impede or undermine their efficacy at the state and local level. However, little to no empirical evidence exists that these assumptions hold. Indeed, non-federal governments often lack the resources and security personnel needed to understand and implement the recommendations supplied in federal security guidance documents. Nor do they have the financial resources to acquire what they lack, or to hire security personnel.

Risk #1. "Best practices" and operational standards make assumptions about the resources available to organizations, and how organizations work, without stating these assumptions explicitly. This leads to an erroneous perception of universal applicability of the "best practices" and standards.

A second risk arises because the guidance is written in a way that makes them hard for non-technical staff to understand. As an example, Federal documents on Internet voting speak of "segmented networks", "S/MIME", and "certificate authorities" [1, 2]. These terms are well known to security practitioners. But many government officials. including network administrators who provide the technical support, are unlikely to understand these terms. If the document were intended for technical support, the people providing that support may not (indeed, generally do not) have the technical, and more specifically the security, expertise needed to implement these recommendations, or determine if in fact the recommendations should be implemented in their environment.

As another example, consider the Center for Internet Security's *FreeBSD Benchmark* [4]. This is a consensus-driven standard developed by volunteers throughout industry and government. While the recommendations are suitable for systems where sharing and access need to be minimized, neither of those characteristics hold in an academic environment, or any environment where sharing resources is a primary goal. The failure to state the goals of the standard in general terms, or explicitly identify the target audience, means that readers who are not technical will not understand the inappropriateness of this standard for other environments.

Risk #2. "Best practices" and operational standards are written for a particular audience, but the *actual* users of that guidance may be a very different audience for whom the guidance is not understood or (worse) misinterpreted.

For instance, federal security guidance normally includes a discussion of SSL (or TLS) and encryption. Applying the SSL (TLS) standards presume the existence of a trustworthy certificate-based public key infrastructure (PKI). But in practice such trustworthy certificate-based management infrastructures rely on certificate authorities that may be vulnerable to attack, as several recent incidents have shown [3].

A more visible problem arises when servers impose requirements on clients. At some point, most users have encountered a web server that recommends viewing the pages in Internet Explorer, or that requires plug-ins that do not work on a large class of systems. A current example is the failure of Google Chrome to support Java 7 on the Mac platform, because of assumptions about word size.[1] Thus, best practices and standards that impose requirements on servers are inapplicable when the clients are incompatible with the server's requirements.

Risk #3. "Best practices" and operational standards make assumptions about the technology and management in which they are implemented and used, often without stating those assumptions explicitly.

In all cases, best practice guides and standards do not exist in a technological and managerial vacuum. Universal optimality is specious, yet guidance documents often are written as if one set of practices will best serve all opera-

[1] Specifically, Google Chrome for MacOS X is a 32-bit web browser, and using Java 7 requires a 64-bit browser.

tional contexts.

We do not know the severity of this problem because of a lack of research and, more fundamentally, a lack of data. These "gaps" in our understanding of cyber security readiness has generated myriad untoward consequences. Take Internet voting as an example. In the United States, laws and regulations containing mandates that specific election activities be available over the Internet assumed substantial state cybersecurity infrastructure. Yet informal discussions with state officers and some public reports have shown these assumptions to be gravely erroneous. As with other essential operations, the election offices of county and state governments differ dramatically, including in their security staffing, knowledge and skill sets, security equipment, security operations (including monitoring and auditing). They may also differ in the scope, frequency, and severity of threats and threat vectors. Relevant differentials range from no wired Internet connectivity to major metropolitan areas whose county or municipal governments maintain multiple internal networks and servers with well-qualified technical staff. Some local governments' technical staff have exceedingly modest technical training and none in security and privacy issues, and work in offices where servers are located in hallways rather than with appropriate physical security. Unquestionably, national efforts to secure cyberspace will be undermined — perhaps profoundly — if this key subset of critical infrastructure continues to lack serious study and solid data.

Understanding whether appropriate levels of security, assurance, and resiliency are achieved requires data that shows how effectively standards and "best practices" are implemented in practice, and data showing that the implementations are appropriate for the environment in which they are implemented.

Risk #4. Lack of data on the *actual practice* of cybersecurity undermines claims of effectiveness of cybersecurity mandates and ability to handle attacks.

The use of standards or "best practice" guidelines requires an understanding of the assumptions underlying the guidelines, the technical and non-technical resources available in the environment in which the guidelines are to be used, and the limitations of the guidelines. If the adoption is mandatory, those who mandate the adoption must also understand these factors, or provide the resources needed to ensure that the standards and guidelines can be implemented correctly, in the target environment, and that those responsible for the implementation understand the guidelines fully.

The ultimate risk is that governments and organizations may lack the expertise to implement standards and best practices effectively, or may apply them in situations where some components should not be applied. The ironic risk looms that a security-motivated organization will rely on omnibus guidance documents that are not well-designed for that organization's context. The organization uses them in an effort to improve their security posture and decrease the likelihood that threats can be realized; yet that reliance, and perhaps faulty implementation, actually increases the likelihood that the threats can be realized.

1. REFERENCES

[1] G. Beier, S. Chokhani, N. Hastings, P. Hoffman, J. Knoke, A. Regenscheid, and S. Shorter. Information system security best practices for UOCAVA-supporting systems. NISTIR 7682, National Institute of Standards and Technology, Sept. 2011.

[2] N. Hastings, R. Peralta, S. Popoveniuc, and A. Regenscheid. Security considerations for remote electronic UOCAVA voting. NISTIR 7770, National Institute of Standards and Technology, Feb. 2011.

[3] N. Leavitt. Internet security under attack: The undermining of digital certificates. *Computer*, 44(12):17–20, Dec. 2011.

[4] T. Rhodes. Freebsd benchmark v1.0.5 (freebsd 4.10 and above). Technical report, The Center for Internet Security, East Greenbush, NY, USA, 2005.