

# **THE CYBER SECURE INSTITUTE'S PRELIMINARY ANALYSIS OF NIST'S RECOMMENDED SECURITY CONTROLS FOR FEDERAL INFORMATION SYSTEMS AND ORGANIZATIONS**

*AUGUST 4, 2009*

The following provides the Cyber Secure Institute's initial and preliminary observations<sup>1</sup> of the National Institute of Standards and Technology (NIST) newly released *Recommended Security Controls for Federal Information Systems and Organizations*<sup>2</sup> (The NIST Recommendations or the Recommendations).

The NIST Recommendations are a critical component of the Federal cybersecurity effort. The Recommendations will shape the security approach of all unclassified Federal IT systems. NIST also sees the final publication of the Recommendations as "represent[ing] a major step toward building a unified information security framework for the entire federal government." Ron Ross, of NIST's computer security division said in the NIST press release that:

"This final publication represents a solidification of the partnership between the Department of Defense, the Intelligence Community, and NIST and their efforts to bring common security solutions to the federal government and its support contractors . . . . The aim is to provide greater protection for federal information systems against cyber attacks."

In addition, how the Recommendations are implemented will have spill over effects on IT security efforts beyond the Federal government, to include both the sub-Federal level public sector and the private sector. And, in turn, they will impact a major portion of the Federal IT market, and the larger IT market as a whole.

---

<sup>1</sup> In the interest of providing timely information to shape the initial discussion of these new NIST Recommendations, this analysis is being provided on an ongoing basis. As a result, our analysis is subject to change on an ongoing, real time basis. Moreover, while this analysis hits key points, it is not intended to be definitive at this time. The Institute expects at some point to revisit the Recommendations and provide a fuller, more in depth analysis.

<sup>2</sup> NIST, Joint Task Force Transformation Initiative, *Recommended Security Controls for Federal Information Systems and Organizations*, NIST Spec. Pub 800-53, Rev. 3, Aug. 2009 (Available at <http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final.pdf>)

## I. Overview

The NIST Recommendations are an important step forward in bringing a more unified, coherent and integrated approach to IT security. They make important security strides in a number of key areas. However, they also raise a number of serious questions, which are addressed below.

## II. Baseline Controls

The NIST Recommendations begin by instructing Federal agencies and organizations to use Baseline security controls as the starting point for their cybersecurity efforts. In short, an agency should first determine if its systems are “Low-Impact,” “Moderate-Impact” or “High-Impact.” This determination drives the specific security control requirements for the applicable system; the NIST Recommendations provide an extensive list of requirements across a range of areas—from lockable casings to metadata flow control—for systems within each classification. The Recommendation’s Baseline levels are based off of the FIPS 199 and 200 standards.<sup>3</sup> FIPS 199 and 200 both provide in key part:

The generalized format for expressing the security category, SC, of an information type is:

SC information type = {(confidentiality, impact), (integrity, impact), (availability, impact)},  
where the acceptable values for potential impact are LOW, MODERATE, HIGH, or NOT APPLICABLE.

They also provide:

The *potential impact* is **LOW** if—

– The loss of confidentiality, integrity, or availability could be expected to have a **limited** adverse effect on organizational operations, organizational assets, or individuals.

AMPLIFICATION: A limited adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is noticeably reduced; (ii) result in minor damage to organizational assets; (iii) result in minor financial loss; or (iv) result in minor harm to individuals.

---

<sup>3</sup> NIST, Federal Information Processing Standards Publication, Standards for Security Categorization of Federal Information and Information Systems, FIPS Pub 199, Feb. 2004 (available at <http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf>); NIST, Federal Information Processing Standards Publication, Minimum Security Requirements for Federal Information and Information Systems, FIPS Pub 200, Mar. 2006 (available at <http://csrc.nist.gov/publications/fips/fips200/FIPS-200-final-march.pdf>).

The *potential impact* is **MODERATE** if—

- The loss of confidentiality, integrity, or availability could be expected to have a **serious** adverse effect on organizational operations, organizational assets, or individuals.

AMPLIFICATION: A serious adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a significant degradation in mission capability to an extent and duration that the organization is able to perform its primary functions, but the effectiveness of the functions is significantly reduced; (ii) result in significant damage to organizational assets; (iii) result in significant financial loss; or (iv) result in significant harm to individuals that does not involve loss of life or serious life threatening injuries.

The *potential impact* is **HIGH** if—

- The loss of confidentiality, integrity, or availability could be expected to have a **severe or catastrophic** adverse effect on organizational operations, organizational assets, or individuals.

AMPLIFICATION: A severe or catastrophic adverse effect means that, for example, the loss of confidentiality, integrity, or availability might: (i) cause a severe degradation in or loss of mission capability to an extent and duration that the organization is not able to perform one or more of its primary functions; (ii) result in major damage to organizational assets; (iii) result in major financial loss; or (iv) result in severe or catastrophic harm to individuals involving loss of life or serious life threatening injuries.

Because the NIST Recommendations are based off of the FIPS 199 and 200 standards it is necessary to first briefly touch on the ramifications of their three-tiered system of impact classification.

While the FIPS 199 and 200 standards provide a thoughtful analytical framework for determining what levels of security are required in a given Federal system, they are not without their own limitations within the context of the NIST Recommendations.

Most importantly, the serious IT security requirements really only kick in at the High Impact Baseline level. At the Moderate and Low levels systems are “not intended to protect a . . . system against high-end threat agents (i.e., threat agents that are highly skilled, highly motivated, and well resourced).” However, so called high end threats are now the norm not the exception. Federal and private sector IT professionals increasingly report that the attacks they confront on a regular basis are from highly skilled, highly motivated and well resourced actors—ranging from the Russian mob, to the Chinese military, to organized cyber-criminals. Additionally, the concern that is gaining greater attention within the IT community is the insider threat, which is almost exclusively made up of highly skilled, highly motivated individuals.

This is particularly troubling if you look at the type of systems that can be classified as Moderate-Impact, and which would remain vulnerable to sophisticated attacks. For example, NIST 199 provides as an example a law enforcement agency managing “extremely sensitive investigative information” would only qualify as a Moderate system. While the Federal perspective may not consider such a system at a risk from “high-end threats,” it seems likely that these sophisticated attackers might have a different opinion.

To summarize this concern, the Baseline Controls provide protections against “highly skilled, highly motivated, and well resourced” threats only for systems designated High Impact. However, the definitional aspects of High Impact systems do not apply to vast numbers of Federal IT systems that could have major impacts on the nation and individual Americans if breached. For example, the e-Health systems now being pushed by the Obama Administration would seem to fall in the Moderate category. However, the threat to so called Low and Moderate Impact systems come from sophisticated actors, like the Chinese military and organized crime. Nevertheless, the NIST recommendations only require these systems to be secure against unsophisticated threats—the proverbial teenage vanity hacker hacking away in the basement.

To be fair, the NIST Recommendations provide that the Baseline determination is the starting point and that the security requirements for a given system then need to be tailored based on specific concerns. However, that does not alter the fact that the more impactful security requirements as a rule only come in to play at the High Impact level, and the bar to be in that category is set rather high.

Additionally, the use of the FIPS 199 classifications raises a concern for individual Americans. From an individualistic perspective, the 199 High Impact classification is based heavily on the loss of life or serious life threatening injury or serious financial loss. An agency’s systems can classify as High Impact, requiring advanced security, if a failure would cause the organization to be unable to perform a core mission or a “major damage to organizational assets . . .” These are important factors but they are inward looking—they reflect solely back on the agency.

Increasingly the Federal government holds critical information about each of us, and its systems have a regular and systemic impact on our day-to-day lives. The Federal IT ramifications for every American are only going to increase given the President’s IT initiatives, such as e-Health. Yet, the only individual impacts that can push a system into the high impact category are death, life-threatening and serious financial harm. These categories are the most critical. However, there are a host of other ways that a Federal IT security failure can have serious impacts on an individual, or for that matter a corporation or company.

Look back at the investigative agency example from FIPS 199. Imagine all the types of information that might be contained in an individual investigative file—

wiretap conversations, surveillance reports, financial records, a person's entire life. If the investigation data breached was, for example, an Securities and Exchange Commission investigation into the actions of a publicly traded company, the breach could have serious impacts on the company, its shareholders, and even the market. Yet, a breach that compromised these types of systems would only seem to measure a Moderate Impact, which in turn drives the level of security across the systems.

Many may brush that aside: "I don't need to worry about those sorts of things. I don't behave in a manner that would make me a target of the FBI or the IRS. I don't need to worry about that." However, that same analysis of the investigative agency would seem at first glance to apply to an e-Health. It seems overwhelmingly likely that the system would qualify only at the Moderate level. However, such an analysis would overlook the serious and far ranging risks that could attach to a breach of the e-Health system. The breach of one person's e-Health database would not cause a major loss of agency's organizational assets and it is hard to say that a lone breach would undermine a core agency function writ large. However, unless the breach put life and limb at risk, no matter what we as individuals might think, the system holding and protecting that information would not rate High-level protection.

All this is to say that the devil in determining the level of security classification is in the details and up to the action agency. And, it is unclear how agencies will react when faced with the choice between levels of security—and the attendant consequences in terms of budgets, personnel, other resources and operational demands.

### **III. Performance Standards**

Perhaps the biggest gap in the whole approach is the general lack of measurable or certified performance standards and validation processes. The NIST Recommendations provide 132 pages of requirements that systems must be able to meet within the three Baseline Control levels.

However, at the end of the day, there is no requirement that technologies or other system elements be tested and proven to meet those requirements. There is no requirement for penetration testing or mathematical proof. There is no third-party validation or verification process or regime. In short, there is no certification requirement.

At the High-Impact Baseline level (systems that we trust with financial calamity and the possible loss of life) the Recommendations do state that developers and implementers must provide "a description of the functional properties and design/implementation of the control with sufficient detail to permit analysis and testing of the control (including functional interfaces among control components)." Additionally, for those systems that must withstand high-end

threats, the Recommendations also provide that a control must be “developed in a manner that supports a high degree of confidence that the control is complete, consistent, and correct.

However, while the developers and implementers are required to provide such information to allow for testing. There is no requirement for the testing, nothing that says how a system should be tested against the requirements it is being deployed to address.

This is surprising because NIST and the National Security Agency are partners in the National Information Assurance Partnership, or NIAP, which has the capacity to test specific products against recognized security profiles and to certify those that meet these standards. If the NIST wanted to impose a certification regime, to actually ensure the level of security across the unclassified Federal IT realm, it could easily have done so by relying on its own partner initiative. However, they declined to take this approach.

As a result, reading the standards, one is left to wonder, who is to say that a system can do all these things? Whose word are we to take? In a marketplace too often replete with grossly exaggerated security claims it seems misguided to put such heavy reliance upon what system developers and implementers tell the Federal government about the security of their own systems.

In fact, the Recommendations make clear at any number of points that they expect that there will be failures. Notwithstanding any requirement to, for example, control access from one domain to a different domain with different security requirements, or ensure contractors meet these requirements, at the end of the day it is still a largely hack and patch approach. Set out a standard, expect failure, look to fix the holes after the fact.

#### **IV. “Trustworthiness”**

This hack and patch approach is particularly clear with respect to what the NIST Recommendations call “trustworthiness.” Trustworthiness is defined as “information systems that are capable of being trusted to operate within defined levels of risk despite environmental disruptions, human errors, and purposeful attacks that are expected to occur in the specified environments of operation.”

One might paraphrase that to read trustworthiness is the ability of a system to remain secure in the real world, against real world threats, doing real world things, being operated by real people.

Interestingly, the NIST Recommendations provide that “typically functions, subsystems, and components are highly integrated, making separation by trustworthiness perhaps problematic and at a minimum, something that requires careful attention in order to achieve practically useful results.” In other words, no

matter how desirable, achieving trustworthiness in the actual Federal IT space would not be easy. To be sure, in the face of today's threats, to actually make the world of Federal IT secure, would be no small challenge. And, it would not come without serious expense. However, at the theoretical level, isn't that the precise point of the NIST Recommendations? Isn't that exactly what President Obama has repeatedly charged the Federal bureaucracy with doing?

Notwithstanding all of the important advances made throughout the NIST Recommendations on the whole, it is the Institute's view that this approach is inherently problematic and threatens to simply perpetuate our reactive and ineffectual current approach to cybersecurity.

## **V. Best Available Cybersecurity Technology (BACT)**

The other major thing not present in the NIST approach is any market- or security- or technology-driver.

In other areas, the Federal government has used standards to drive technology to new levels. For example, a host of environmental laws include "Best Available Control Technology" requirements. Such standards create a market for any technology that can prove that it produces a measurable benefit over and above the best prior technology.

Similarly, the Federal Government has also used its buying power to drive change. In these areas, the Federal government has essentially said, "if you want to sell to us, you need to go beyond the private sector requirements and push the envelope of change to drive policy and behaviors." For example, the Clinton Administration's Executive Order 12873 drove a Federal market for recycled content paper. Today, 98 percent of all paper purchased by the Federal government contain 30 percent recycled content.

NIST could have recommended that, where applicable, within each of these delineated cybersecurity requirements, Federal agencies should use the "Best Available Cybersecurity Technology." The strength of "best available" requirements on the whole is that they create a constant market pressure for better, more robust, higher trustworthiness technologies and systems. Whoever, pushes the technological envelope instantly becomes the market default—until such time as they too get surpassed.

Such an approach would have harnessed the Federal IT market, which is substantial, to drive the larger market for truly secure cybersecurity technologies. Moreover, it would have used the Federal IT world to help create the de facto standard for cybersecurity more broadly.

However, NIST did not take this approach. (In fact, without suggesting any value judgment, there is no way to know if such an approach even was considered within the NIST effort here.)