Content type: Orig-research
True
NIST Contributions to IT

NIST and Computer Security

William Burr, Hildegard Ferraiolo, and David Waltermire, US National Institute of Standards and Technology

The US National Institute of Standards and Technology's highly visible work in four key areas—cryptographic standards, role-based access control, identification card standards, and security automation—has and continues to shape computer and information security at both national and global levels.

The US National Institute of Standards and Technology (NIST) touches many IT areas; one of the most visible and well known is computer and information security. In many ways, NIST (known prior to 1988 as the National Bureau of Standards, or NBS) has shaped the computer security field since its early days and continues to do so today.

Here, we review four key areas—cryptographic standards, role-based access control (RBAC), identification card standards, and security automation—that illustrate both NIST's impact and the variety of ways in which NIST interacts with industry. All four areas are widespread components of the international information infrastructure. For example, NIST cryptographic standards are used nearly everywhere for banking and electronic commerce, and have been built into operating systems and communication protocols for decades. NIST also conducts original research and development of technologies that later become industry standards. RBAC is one such example; much of the early research was done at NIST, laying the foundation for academic research and IT industry implementation. More recent work in identity cards and security automation shows how NIST standards have become the basis for products in the IT industry.

Cryptographic Standards

In 1883, Auguste Kerckhoffs formalized the basic idea of encryption standards: the security of encrypted data (ciphertext) rests only upon the secrecy of an encryption key, and not upon keeping the encryption algorithm a secret.[1,2] Still, an encryption scheme must be secure when an adversary knows how it works. How is it possible to have a standard for broad commercial use but keep the algorithm a secret? Strong encryption remained the exclusive preserve of secret intelligence organizations (that did try to keep algorithms secret) until 1973, when NBS solicited proposals for a Data Encryption Standard (DES). IBM submitted the winning proposal (the only complete proposal), which encrypted 64-bit data blocks under a 56-bit key and was adopted as Federal Information Processing Standard (FIPS) 46 in 1978.[3]

Data Encryption Standard

DES was the first example of an open standard for strong cryptography. It helped stimulate cryptography as a field of study in computer science, as well as furthering the development of a worldwide cryptographic research community.

DES was originally motivated by the need to secure the communications of Automated Teller Machine (ATM) networks, and it was soon used widely in the financial services industry. With the Internet's arrival in the 1990s, it was clear that cryptography was vital to securing the net as a basic instrument of commerce. By the late 1990s, however, it was apparent that the DES had a serious shortcoming: the 56-bit key was no longer big enough to protect against attacks by powerful computer networks or special purpose cracking machines.

Advanced Encryption Standard

In 1997, NIST called for candidate algorithms for an Advanced Encryption Standard (AES) that would support keys of 128-, 192-, and 256-bits to encrypt 128-bit data blocks. To encourage participation from the international cryptographic community, NIST held three open AES candidate conferences immediately following or preceding major cryptographic conferences.

The 15 initial candidates were winnowed to five finalists for more intense study. Many papers were published on the candidates, their cryptanalysis, and their performance. NIST selected the Rijndael algorithm—designed by Belgian cryptographers Vincent Rijmen and Joan Daemen—to be the AES and published FIPS 196 in 2001. Today, AES has largely supplanted DES, and many modern computers include special instructions to accelerate it.

Secure Hash Algorithm

In practical terms, public key digital signature algorithms require too much computation to be directly applied to long messages. Thus, hash functions produce a small "message digest" from a long message and are used for many cryptographic functions—most importantly as proxies for the actual messages with digital signatures. For this application, hash functions must be "collision resistant"—that is, it must be computationally infeasible to find two messages, $M_1 \neq M_2$, such that Hash(M_1) = Hash(M_2). NIST has standardized the SHA family of hash functions, which includes SHA-1 (a 160-bit message digest) and SHA-2 (which produces 224-, 256-, 384-, and 512-bit digests).

In 2004 and 2005, a wave of new collision attacks began casting doubt on the security of most existing widely used hash functions, including SHA-1, and caused considerable concern in the cryptographic community. In 2007, NIST responded by calling for a SHA-3 hash function competition to provide a new hash function standard, intended to be very different in operation from SHA-2. Participation was unprecedented: NIST got 64 complete and proper submissions from all over the world.

Three SHA-3 conferences were held in conjunction with major cryptographic research conferences to winnow the candidates to five finalists and intensively study those finalists. Hundreds of papers were published and presented in many international venues. Several major independent websites were set up to support the competition. A huge amount was learned about hash function design. Researchers fabricated several custom application-specific integrated circuits (ASICs) of all five finalists and measured their performance.

The winning algorithm, Keccak—developed by Guido Bertoni, Joan Daemen (again), Michaël Peeters, and Gilles Van Assche—is a major departure from previous hash function designs, with a whole new general design for hash functions: the "sponge construction." It's remarkably flexible and easily adjusted for tradeoffs between speed and security level. Unlike most previous hash functions, Keccak is based on a fixed permutation. Keccak is very fast in hardware, and a little faster than SHA-2 in most laptop, desktop and server computers. The algorithm will also offer variable length outputs and can provide an authenticated encryption service as well.

It will take some time to deploy SHA-3, and many years to properly study and exploit all of the possibilities of the sponge construction and the Keccak permutation. Still, after years of research, no attack on the full SHA-2 has been found. NIST is confident that SHA-2, which is available now, is secure and will serve for a long time.

NIST's International Impact

Although NIST's purview extends only to US federal agencies, as the first pioneer of open cryptographic standards, its encryption and hash standards are among the most widely used algorithms in the world and are foundational for Internet security. The cryptographic competitions are a challenge for NIST's small cryptographic group and aren't warranted for all cryptographic standards and recommendations; however, they can effectively harness the immense energy and collective expertise of the global cryptographic research community, and they help attract the greatest possible scrutiny and cryptanalysis for proposed fundamental standards. Most importantly, when people everywhere trust one of these algorithms, they're relying not just on NIST but also on a vast community of researchers that participate in the NIST process to vet the selection.

Role-Based Access Control

A 2011 economic impact study by the Research Triangle Institute found that most organizations with more than 500 employees use some form of RBAC, which is implemented in many prominent products such as the Microsoft Exchange Server. Yet less than 20 years ago, computer security was most commonly implemented with only access control lists (ACLs) based on user IDs and groups. NIST's contributions in this area were to develop the first general RBAC model, which resolved weaknesses in the ACL approach, added features to simplify privilege administration, and established a formal RBAC standard.[4]

Roles and Access: A Brief History

Roles with different privileges and responsibilities have long been recognized in business organizations, and commercial computer applications dating back to the 1970s implemented limited forms of access constraints based on the user's organizational role. For example, online banking applications in that period included both teller and teller supervisor roles that could execute different sets of transactions, while ATM users were able to simultaneously execute another set of transactions against the same databases.

In the late 1980s and early 1990s, researchers began recognizing the virtues of roles as an abstraction for managing privileges within applications and database management systems. A role was seen as a job or position within an organization. These role-based systems were relatively simple and application-specific, and they evolved on an ad hoc basis within various organizations. That is, there was no general-purpose model defining how access control could be based on roles, and little formal analysis of the security of such systems existed.

In 1992, NIST studied commercial and government organizations and found that access control needs at that time weren't being met by the available commercial products, many of which implemented only discretionary controls and access control lists. In many enterprises in industry and civilian government, users don't "own" the information that they can access. For these organizations, the corporation or agency is the actual owner of system objects, so discretionary control on the user's part might not be appropriate. Conventional multilevel security with military-style classification levels is also inadequate for these organizations, where access is based primarily on job function.

Enter the RBAC Model

NIST proposed a solution to meet these needs in 1992.[5] The solution integrated features of existing application-specific approaches into a generalized RBAC model. The proposal formally described the sets, relations, and mappings used to define roles and role hierarchies, subject-role activation, and subject-object mediation, as well as the constraints on user-role membership and role-set activation. Three basic rules were required:

- *Role assignment.* A subject can execute a transaction only if the subject has selected, or been assigned to, a role. Thus, all active users are required to have some active role.
- *Role authorization.* A subject's active role must be authorized for the subject. With rule 1, this ensures that users can take on only roles for which they are authorized.
- *Transaction authorization.* A subject can execute a transaction only if the transaction is authorized for the subject's active role. Combined with rules 1 and 2, this ensures that users can execute only transactions for which they are authorized.

A key feature of this RBAC model is that all access is through roles. A role is a collection of permissions, and all users receive permissions only through their assigned roles. Within an organization, roles are relatively stable, while users and permissions are numerous and might change rapidly. A superficial similarity exists between roles and groups, but a group is normally implemented as a collection of users, rather than a collection of permissions, and permissions can be associated with both users and the groups to which they belong.

Because users can access objects based on either their user or group ID, it's possible for them to retain access permissions that should be revoked when group permission is removed from the object. RBAC's requirement of access only through roles strengthens security by eliminating this loophole. Additional features of this model are that roles are hierarchical (they can inherit permissions from other roles) and that

they include provision for constraints, including separation of duty.

RBAC Evolves

Under a NIST Small Business Innovative Research competition, SETA Corporation and Ravi Sandhu of George Mason University developed an RBAC family of models. In 1996, Sandhu and his colleagues[6] further broadened the field by introducing *RBAC96*, a framework that breaks down RBAC into four conceptual models.

The base model, RBAC0, contains the minimal features of a system implementing RBAC. Two advanced models, RBAC1 and RBAC2, include RBAC0 and add support for hierarchies (in RBAC1) and support for constraints such as separation of duty (in RBAC2). A fourth component, RBAC3, includes all aspects of the lower-level models. The RBAC96 framework established a modular structure for RBAC systems, offering simplified commercial implementations that provide basic RBAC0 functionality or more advanced features, depending on customer requirements.

In 2000, NIST initiated an effort to establish an international consensus standard for RBAC, publishing a proposal in the ACM RBAC workshop.[7] The proposed standard, known as *the NIST model*, unified the 1992 NIST RBAC definition with the structure of RBAC96. A later revision incorporated features developed through subsequent discussions and formal comments from the research and commercial vendor communities.[8] In 2002, the revised proposed standard was submitted to the international standards process; at that point, commercial firms had already begun building products that conformed to the NIST model.

What is most striking about RBAC's history is its rapid evolution from a concept to a deployed commercial implementation. RBAC differs from many other security concepts in that its costs of deployment need not be justified based solely on perceived threats and system vulnerabilities. Although RBAC allows for the enforcement of a wide variety of important access control policies that are either impractical or even impossible to enforce in its absence, RBAC's productivity advantages alone are often sufficient to justify its deployment. RBAC emerged as the primary enterprise access control model because it was much better suited to commercial users' needs compared to earlier models. (A more detailed discussion of RBAC's evolution is available elsewhere.[4])

Identity Card Standards

Since the Brooks Act in 1965, the US Congress has turned to NIST to develop IT security standards that are both practical and effective—a combination that is often difficult to achieve. Done right, these standards often become the basis for commercial products used by millions. One such example is the Personal Identity Verification (PIV) card, now a key component for both information system and physical access control.

Initiated in 2004 by Homeland Security Presidential Directive 12, the PIV card deployment was aimed at eliminating wide variations in the quality and security of authentication mechanisms used across federal agencies. The mandate called for a common identification standard to promote interoperable authentication mechanisms at graduated security levels based on the environment and data sensitivity. In response, NIST published the 2005 FIPS 201, which specified a common set of identity credentials in a smart card form factor—the PIV card—which is used today government-wide, as intended, for both physical access to government facilities and logical access to federal information systems.[9]

The 2005 release of FIPS 201 marked the beginning of a learn-design-develop-test-validate phase for both the private sector and federal departments and agencies. By 2009, more than 300 standard-conformant products had been developed, validated, and brought to market in support of the PIV card and its infrastructure. Departments and agencies also developed and refined their PIV card issuance processes. PIV card issuance systems are have been operating, and close to 5 million PIV cards have been issued to federal employees and contractors, according to the Office of Management and Budget (OMB).[10] Today, the emphasis has shifted from PIV card issuance to its deployment and use in logical and physical access applications. Many applications now use the PIV card for government network access, including MyPay, Employee Express, and OMB Max Portal.

Security Automation and Vulnerability Management

Most organizations face tough questions related to the IT infrastructure that supports critical business and mission objectives. What software do I use in my systems? Is that software vulnerable to attack? Is it properly configured to reduce my attack surface? The NIST Security Automation Program developed reference data and technical specifications for sharing security information between information systems that lets commercially available solutions answer these questions.

Security automation specifications and reference data can be used to maintain enterprise system security, including detecting the presence of installed software, automatically verifying the installation of patches, checking systems security configuration settings, and examining systems for indicators of compromise. In addition to configuration, patch, and compliance use cases, structured data using security automation specifications are also being leveraged to solve problems related to software assurance, asset inventory, malware detection, event correlation, and continuous monitoring. Through the creation of flexible, open standards and international standards recognition, security automation aids in IT infrastructure interoperability, broad acceptance, and adoption, and helps create opportunities for innovation.

Security Content Automation Protocol

Supporting the overarching security-automation vision requires both trusted information and a standardized means for sharing it. Through close work with its government, academic, and industry partners, NIST developed the Security Content Automation Protocol (SCAP)[11] to support information assurance by providing the standardized data formats needed to share information between endpoint devices and enterprise components that aggregate, store, and analyze the data. SCAP provides an automated means for collecting security and operationally relevant data.

SCAP overview. The SCAP suite of specifications uses eXtensible Markup Language (XML) to standardize how security software products communicate information about the endpoint state. SCAP is a multipurpose protocol that supports automated software inventory, configuration and vulnerability checking, security control compliance activities, security measurement, and the identification of malware and compromise indicators.

SCAP consists of

- standardized identifiers for software names, configuration items, and vulnerabilities;
- data formats that enable content-driven data collection of the endpoint state, the evaluation and reporting of collected data, and characterization of device identities;
- scoring methodologies for measuring the relative security impact of software flaws and misconfigurations; and
- guidance on using XML digital signatures and cryptographic standards for protecting the integrity of content used for data collection and the data reported as a result.

In September 2012, NIST Special Publication (SP) 800-126 revision 2 was released, providing the specification for SCAP version 1.2, which is the most current SCAP version.

Penetration and validation. The US government, in cooperation with academia and private industry, is adopting SCAP and encourages its use to support security automation activities and initiatives. This adoption has made SCAP a significant component of information security management and governance programs in government and industry. To promote the ongoing adoption and maintenance of SCAP, NIST operates the SCAP Validation Program, which conducts formal conformance testing through third-party testing laboratories. The SCAP Validation Program ensures that products correctly implement SCAP as defined in SP 800-126. Conformance testing is necessary because SCAP is a complex specification consisting of 11 individual specifications that work together to meet various use cases. A single error in product implementation could result in undetected vulnerabilities or policy noncompliance within agency and industry networks.

The SCAP Validation Program was created in 2008 to support several federal configuration

standardization initiatives. The program coordinates its work with the NIST National Voluntary Laboratory Accreditation Program to establish independent conformance testing laboratories that test based on the SCAP Validation Program Test Requirements. When testing is completed, the laboratory submits a test report to NIST for review and approval. SCAP validation testing has been designed to be inexpensive and effective. The SCAP conformance tests are either easily human-verifiable or automated through NIST-provided reference data and tools. As of July 2013, the program has eight accredited independent laboratories and has validated 50 products from 33 different vendors. The SCAP Validation Program is expanding to provide enhanced testing support and will evolve to include new technologies as SCAP matures. Expansion plans include support for additional federal configuration baselines, expanded SCAP validation test content, and expanded automated testing capabilities.

National Vulnerability Database

Established in 1999, the National Vulnerability Database (NVD) is the US government repository of security automation reference data, which is based on security automation specifications. The NVD reference data provides a standards-based foundation supporting the automation and measurement of software assets, vulnerability, and security-configuration management; security measurement; and compliance activities. NVD data is a fundamental component of NIST's security automation infrastructure and is substantially increasing network security worldwide.

The NVD provides structured information on more than 57,000 software vulnerabilities. Each vulnerability is associated with a Common Vulnerabilities and Exposure Identifier (one of the SCAP identification formats). Vulnerability data provided by the NVD is used to score each vulnerability record using the Common Vulnerability Scoring System version 2, which lets users compare the relative impact of vulnerabilities. The NVD vulnerability data also contains references to government, vendor, and third-party advisories, and links to patches (when available). These references include links to more than 240 US Computer Emergency Readiness Team (US-CERT) alerts and more than 2,700 US-CERT vulnerability summaries.

The NVD vulnerability data also includes mappings to vulnerable products, categorizations of the vulnerability type using the Common Weakness Enumeration (CWE), and references to more than 8,100 automated vulnerability assessment checks that can be used by SCAP validated tools to detect vulnerabilities on an endpoint. Using this information, IT security tool vendors can augment their data feeds and services; security researchers can analyze various aspects of vulnerabilities, including trends; and user organizations can help analyze vulnerability management activities.

Finally, the NVD maintains a catalog of more than 220 security configuration checklists for IT products provided by government, product vendors, and third-party organizations. These checklist entries point to prose guides, configuration scripts, and automation content used to assess IT product configurations. More than 50 of these entries point to configuration guides defined using SCAP; with these, users can automatically check recommended configuration settings using SCAP-validated tools. The use of NVD SCAP data by commercial security products deployed in thousands of organizations worldwide has extended NVD's effective reach.

With a mission to promote US innovation and industrial competitiveness, NIST requires a unique set of capabilities in both research and standards-setting In the area of cryptography, NIST is now in the process of specifying a SHA-3 standard that will lead the way a new complete family of standardized symmetric key functions that are based on the permutation that underlies KECCAK, allowing the full range of symmetric key functionality (such as an authenticated encryption and pseudorandom number generation) to be generated from one simple, efficient primitive. For security automation, NIST continues to work with industry and academia to advance and promulgate consensus-based, international standards that support commercially available solutions to address software inventory, configuration management, and indicator sharing applications that use security automation data and techniques. Finally, identity management card standards has broadened its focus at NIST by finding adaption in mobile devices authentication and cloud services.

Some commercial entities, equipment, and materials are identified in this article to describe an experimental procedure or concept; such identification isn't intended to imply recommendation or endorsement by the US National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

References

1. A. Kerckhoffs, “La Cryptographie Militaire,” *J. Des Sciences Militaires*, vol. IX, Jan. 1883, pp. 5–83.
2. A. Kerckhoffs, “La Cryptographie Militaire,” *J. Des Sciences Militaires*, vol. IX, Feb. 1883, pp. 161–191.
3. *Data Encryption Standard*, *FIPS-Pub. 46*, Nat’l Bureau of Standards, US Dept. Commerce, Jan. 1977.
4. V.N.L. Franqueira and R.J. Wieringa, “Role-Based Access Control in Retrospect,” *Computer*, vol. 45, no. 6, 2012, pp. 81–88.
5. D.F. Ferraiolo and D.R. Kuhn, “Role Based Access Control,” *Proc. 15th Nat’l Computer Security Conf.*, 1992, pp. 554–563; www.csrc.nist.gov/rbac/ferraiolo-kuhn-92.pdf.
6. R.S. Sandhu et al., “Role-Based Access Control Models,” *Computer*, vol. 29, no. 2, 1996, pp. 38–47.
7. R. Sandhu, D. Ferraiolo, and R. Kuhn, “The NIST Model for Role-Based Access Control: Towards a Unified Standard,” *Proc. 5th ACM Workshop on Role-Based Access Control*, ACM, 2000, pp. 47–63; http://csrc.nist.gov/staff/Kuhn/towards-std.pdf.
8. D.F. Ferraiolo et al., “Proposed NIST Standard for Role-Based Access Control,” *ACM Trans. Information and System Security*, vol. 4, no. 3, 2001, pp. 224–274.
9. *Personal Identity Verification of Federal Employees and Contractors*, *Federal Information Processing Standard 201-2*, US Nat’l Inst. Standards and Technology, Sept. 2013.
10. “HSPD-12 Public Report Summary,” White House, 2013; www.whitehouse.gov/sites/default/files/omb/assets/egov_docs/hspd-12_reporting_workbook_q2fy2013_public_report.pdf.
11. D. Waltermire et al., *The Technical Specification for the Security Content Automation Protocol (SCAP)*: *SCAP Version 1.2*, special publication 800, US Nat’l Inst. Standards and Technology, 2011, p. 126.

Bill Burr *managed the NIST Cryptographic Technology Group from 2000 through 2010. He organized and led the SHA-3 competition through the selection of five “finalist” algorithms. Bill joined NIST in 1978 to work on computer peripheral interface and high speed local area network standards, and since 1990 he has worked in the Computer Security Division on standards for PKI and cryptography. Bill was the lead author of the influential NIST Special Publication 800-63 which gives technical guidance on authentication to Federal agencies implementing Internet E-Government services. Bill chaired the Federal Public Key Infrastructure Technical Working Group for about a decade and led the final selection round of the Advanced Encryption Standard (AES). Contact him at William.burr@nist.gov*

Hildegard Ferraiolo *is a Computer Scientist at the US National Institute of Standards and Technology. Her research interests include Identity Management, Mobile Device Security and Authentication, Smart Cards, Biometrics, Public Key Infrastructure (PKI), Cryptography. She is the program manager for the HSPD-12/PIV program. Her recent focus is to expand the smart card-based identity scheme to mobile devices and cloud services. She has received the 2007 Department of Commerce Gold Medal. Contact her at hferraio@nist.gov.*

David Waltermire *is the Specification Architect for the Security Automation Program at the US National Institute of Standards and Technology. His research interests include standards-based, platform neutral techniques for automating software inventory and configuration management practices, architectural design patterns for monitoring the use and effectiveness of security controls, and endpoint management protocol design. Waltermire has authored/co-authored a number of specifications in the areas of security automation and continuous monitoring of security controls. Contact him at david.waltermire@nist.gov.*

The US National Institute of Standards and Technology's highly visible work in four key areas—cryptographic standards, role-based access control, identification card standards, and security automation—has and continues to shape computer and information security at both national and global levels.

Keywords: authentication, e-commerce, identity, smart card, cryptographic competition, hash function, Security Content Automation, mobile device authentication