Emerging Standards

Editors: Ramaswamy Chadramouli, mouli@nist.gov
Rick Kuhn, kuhn@nist.gov
Susan Landau, susan.landau@sun.com

Security and Privacy Landscape in Emerging Technologies

Susan Landau
Sun Microsystems

This column has run for three years, and, as with all good things, it's time for it to come to an end. For its final installment, I'll focus on some just emerging standards and technologies, as well as present some near-term standards challenges.

Just emerging standards

Meeting needs at the intersection of security and safety can be especially challenging, particularly when standardized interfaces are also required. Critical infrastructure systems that were once isolated are now rapidly being interconnected, often while engineers are simultaneously introducing significantly stronger security measures than before. Fortunately, several organizations are making progress on standards for security in critical infrastructure systems. Two of the most interesting domains are industrial control systems and emergency management, which are both characterized by severe requirements for both security and safety.

Industrial control systems

Among the fastest growing needs in IT security is better protection for industrial control systems (ICSs), including supervisory control and data acquisition (SCADA). These systems help monitor and control large infrastructure frameworks (including electrical power generation, transmission and distribution, oil and gas transport, and water pumping and supply systems). They often span large geographic areas and rely on various communication systems. ICS applications are unique in their architecture in that they might comprise both embedded (programmable logic controllers) and conventional computing devices and applications. ICSs were originally closed-loop systems with low security exposure. More recently, factors such as market pressure (requiring that control systems be connected to planning systems such as enterprise requirements planning) and deregulation (necessitating a need for energy generators or distributors to give energy traders access to SCADA systems) have altered these systems' connectivity topology, making them more widely accessible and thus increasing their security exposure.

Integrating security into a domain that was previously open is a familiar problem for security practitioners. However, many solutions that made sense for conventional IT systems might not work with ICSs, which often have very different performance requirements. A slow authentication step that's annoying in a point-of-sale system could have drastic consequences in a power-plant system that requires prompt action from operators. Given ICS applications' unique architecture and connectivity topology, ICS designers must develop security metrics, solutions, and test methodologies. Standards that address the special needs of ICS security are already appearing but aren't as mature as those in other domains. They include ISA-99 and the US National Institute of Standards and Technology (NIST) SP800-53 publication with supplemental guidance for ICS (see http://csrc.nist.gov/publications/nistpubs/800-53-Rev2/sp800-53-rev2-final.pdf). Developed by the international Instrumentation, Systems, and Automation Society (ISA), ISA-99 addresses security for safety or infrastructure-critical systems (see www.isa.org/MSTemplate.cfm?MicrositeID=988&CommitteeID=6821). The ISA is still developing parts of ISA-99, which might be complete in 2009 or 2010. NIST SP800-53 is a comprehensive standard for US government agency security, recently supplemented with recommendations for ICS security. While ICS security planners await completion on the international standard, they might benefit from NIST recommendations. The ISA is mapping some of ISA-99's foundational requirements to the controls specified in SP 800-53, and parts of ISA-99 use a format similar to SP 800-53.

Emergency management systems

In national emergency situations (both man-made and natural), federal, state, local, tribal, and nongovernmental organizations that provide emergency response and management services need to securely exchange data and share information. Emergency management or incident management systems exhibit considerable heterogeneity across organizations and functional disciplines (fire departments, emergency medical services, and so on), so the key to improving emergency response is an interoperable, platform-neutral standardized message framework. To fill this need, the Organization for the Advancement of Structured Information Standards' (Oasis's) Emergency Management Technical Committee, in cooperation with the US Department of Homeland Security/Federal Emergency Management Agency (as part of the Disaster Management eGov Initiative) and the Emergency Interoperability Consortium initiated the Emergency Data Exchange Language (EDXL) activity.

According to the memorandum of understanding between these three entities, the EDXL was proposed as a "cooperative effort to define a NIMS [National Incident Management System]-compliant family of shared data-exchange specifications encompassing incident notification and situation reports; status reporting; resource requests and dispatch; analytical data; geospatial information; and identification and authentication."

The first specification developed was EDXL Distribution Element (EDXL-DE) version 1.0, which provides standardized routing assertions for all types of emergency data including XML messages, spreadsheets, JPEG images, or other digital data. In other words, we can think of the "distribution element" as a container that provides information needed to route payload message sets (such as alerts or resource messages) by including key routing information such as distribution type, geography, incident, and sender/recipient IDs.

The next logical step was to standardize payload message sets. As mentioned earlier, alerts and resources form two key message categories. Accordingly, the Oasis Emergency Management Technical Committee began work on a specification that contains a suite of standard messages for data sharing among emergency and other information systems that deal in requesting and providing emergency equipment, supplies, people, and teams. This specification, the EDXL Resource Message (EDXL-RM), defines 16 separate message types supporting the major communication requirements for resource allocation across the emergency incident life cycle. This includes preparedness, prestaging of resources, initial and ongoing response, recovery, and demobilization/release of resources. The committee published the third public draft of EDXL-RM version 1.0 on 27 May 2008. Apart from alerts and resources (in terms of equipment), medical facilities' (or hospitals') availability is key to an emergency response and recovery effort. Thus standards are needed for different emergency management jurisdictions to exchange data about hospital facilities in an interoperable way. To address this need, the Oasis Emergency Management Technical Committee has developed and published the document titled, "Emergency Data Exchange Language Hospital Availability Exchange (HAVE) Version 1.0, Public Review Draft 03" (http://docs.oasis-open.org/emergency/edxl-have/pr03/emergency_edxl_have-1.0-spec-pr03.html). EXDL-HAVE describes a standard message for data sharing among emergency information systems using the XML-based EDXL with regard to facility availability for treating victims of various types of incidents.

Emerging challenges

In addition to these newly emerging standards, I consider some technologies whose privacy and security standards have yet to fully emerge. This lack is unsurprising; social changes often follow technology development. Consider the automobile's effect on US housing patterns, or the cell phone on young people's patterns and behaviors worldwide. As human behaviors change, new technologies arise that accommodate the new social patterns. This is a complex dance, and the technological standards required to fully enable the new activities sometimes lag behind, rather than lead. Such has been the case with communication privacy and security standards.

In 1994, Netscape launched the first Web browser and SSL 1.0, and e-commerce became possible. It was the beginning of the Web as well as the start of the privacy and security threats that are now daily news (and are much of the reason for this magazine). The year 2000 marked the dot-com meltdown, and the phoenixes that arose from these ashes—Web 2.0 and globalization—bring numerous security and privacy challenges.

Medical facilities' availablity is key to an emergecy response effort. Thus standards are needed for different emergency management jurisdictions to exchange data about hospital facilities in an interoperable way.

Challenges in Web 2.0

In this decade, we've seen Web 2.0 blossom, a growth that encompasses both the Web as a service provider and greater Web interactivity—as compared to the situation pre-2000—engendered by newer browsers and sites. This column hasn't focused on emerging privacy standards in Web 2.0 applications because such standards don't really exist. Some fundamental work must first occur.

Yes, various technologies enable us to securely exchange data, including Transport Layer Security (TLS), SSL, and Secure Socket Shell, but the basic issue of how users convey their privacy preferences remains unsettled. Privacy technology is really about access control, control in "the other direction," as it were. What happens when a user stores private information on someone else's server? How can the user control the way the provider uses that data?

Consider, for example, how Google's calendar server protects the privacy of users' calendar data. The Google calendar's privacy pages state that usage data is preserved for 90 days in aggregate form, but no set of standards exists—a user or organization can't simply say, "I want to know how the data is aggregated, how many users' data is aggregated, and what, if anything, can be discovered about me," or "I will permit this site to include me in only this type of data aggregation (but not that)." Similarly, although Microsoft's HealthVault—a service that lets users manage their healthcare information—has a privacy policy, users might inadvertently let someone remove control of their own health data (by giving custodial access to someone who might remove users' custodial status, for example). No set of standards exists across Web sites that lets a user say, "No one shall have the right to remove my control of the data I entered into this site," or "The data I provide here can be used by the site but only if it's aggregated in groups of at least 5,000."

Privacy technology is really about access control, control in "the other direction." What happens when a user stores private information on someone else's server? How can the user control the way the provider uses that data?

This area is ripe for research, development, and standards establishment. The problems are extremely important, but are also difficult ones, and thus so far, there has been only preliminary work tackling the complex problem of translating policy statements, such as the Fair Information Practices, into anything resembling code. Nothing approaching protocols or standards yet exists.

Web 2.0's other challenging security and privacy aspect is *identity management systems*, which *S&P* covered extensively in the March/April issue (vol. 6, no. 2, pp. 13–57). The Security Assertion Markup Language (SAML, which underlies the Shibboleth and Liberty efforts[1]) and InfoCard have developed single sign-on (SSO) standards with some privacy and security protections—SAML has been more heavily architected for such protections—whereas OpenID, a lightweight SSO, eschewed privacy and security protection in its initial specifications. Users need privacy and security protections, whereas developers want to balance such protection with ease of implementation. The latter is a complex challenge, and new standards might yet emerge in this domain.

Globalization factors in standards adoption

In part because of the fiber buildout of the 1990s, globalization of services occurred remarkably quickly after the dot-com meltdown. Other globalization changes happened as well. China, having embraced a form of capitalism, developed as an economic powerhouse and stepped into the computing industry, including purchasing IBM's personal computer business, which was renamed Lenovo.

China also began to exercise its own national efforts in cryptographic standardization. To understand the effects of this, let's look at recent bit of history in developing cryptographic standards.

In 1973, the US National Bureau of Standards issued a call for submissions for a block cipher; IBM responded with the Feistel cipher on a 64-bit block using a 56-bit key (in a Feistel cipher, a $2t$-bit input is split into two halves; in the i^{th} round of the cipher, the right half from the previous round becomes the new left half, while the new right half is XORed with the previous left half and a function of the round subkey). The IBM submission eventually became the Data Encryption Standard (DES), but the algorithm was controversial throughout its 30-year history because some doubted its security. Concerns existed that the key length was too short (although DES only fell to a brute-force attack in 1998) and that the US secret cryptanalytic organization, the National Security Agency (NSA), had a secret backdoor into the algorithm that would let it crack encrypted messages quickly.

In the late 1990s, as DES's vulnerability to a brute-force attack became clear, NIST (formerly the National Bureau of Standards) issued a call for a new algorithm. The selection process was highly transparent this time. NIST ran

an international competition, and the Advanced Encryption Standard (AES), approved in 2001, was the work not of Americans, but of two Belgian cryptographers.[2] The transparency of the selection process fostered AES's immediate acceptance in almost every part of the world, an extremely useful outcome if you're interested in deploying secure systems.

In one, country, however, acceptance of AES was not so immediate—China. The conflict started indirectly. In 2003, a semiconductor spin-off from Xiadin University proposed a proprietary block cipher, WAPI, as an ISO 802.11i standard. The Chinese government indicated that it might require WAPI for all Wi-Fi systems sold in China even while the government would restrict access to the technology to Chinese companies.[3] The algorithm itself wasn't made public.

Unsurprisingly, the ISO didn't approve WAPI as an 802.11i standard, but then several interesting things happened.[3] The algorithm, an unbalanced Feistel cipher, was released, and the Chinese government required that all Trusted Platform Module (TPM) deployments in China use WAPI, renamed SMS4, instead of AES as the block cipher.[4]

For the past several years, an industry partnership—the Trusted Computing Group (TCG; www.trustedcomputinggroup.org)—has been developing open specifications for trusted computing. The TPM is the hardware support underlying these specifications; it stores keys, passwords, and digital certificates. TPM chips have been shipping in laptops and PCs for several years, and support is now being provided for them in servers as well. The Chinese government began requiring that all TCG-enabled machines in China use a "TCM," which differs from the TPM in three ways:

- instead of using RSA for public-key cryptography, the TCM uses the Elliptic Curve Cryptosystem;
- the TCM uses Chinese hash algorithms; and
- the TCM provides support for the block cipher SMS4, rather than the more obvious choice, AES. (For export control reasons, TPMs don't support AES encryption of arbitrary data.)

Given China's large population and clear economic clout, it's not surprising that TCG didn't insist that China use the standard TPM. But including SMS4, an algorithm that hasn't been subjected to broad public scrutiny, is an unfortunate step in terms of the transparent process that the open vetting of AES—and thus its near global acceptance—had started. Although other nations have similarly sought to substitute their home-grown solutions for internationally accepted security standards, they couldn't apply the same force to bear on the issue as China could, and therefore didn't succeed. As we go to press, TCG is proposing TPM as an ISO standard. Its approval would strengthen any possible WTO case against the Chinese government's requirement that TCG-enabled equipment use TCMs in China.

It seems likely that China will try similar security substitutions in the future. Security is easiest to achieve when complexity is minimized, when implementation is simple—and when there aren't competing algorithms or APIs with competing implementations. Different security solutions in different nations will fragment security standards, which would, in turn, lead to weakening cybersecurity, a situation no one wants. It isn't clear how this issue will be resolved.

Over the next several years, the technical challenges Web 2.0 provides and the policy challenges globalization generates will undoubtedly provide a quite complex security and privacy standards landscape. □

Acknowledgments

I'd like to express my strong thanks to Ramaswamy Chandramouli and Rick Kuhn of NIST for their clear explanations of standardization efforts in industrial control and emergency management systems. This is our last column as coeditors. Putting together a series of columns on emerging standards and technologies in security and privacy can present its own set of challenges, and my co-editors have been great.

References

1. E. Maler and D. Reed, "The Venn of Identity: Option and Issues in Federated Identity Management," *IEEE Security & Privacy*, vol. 6, no. 2, 2008, pp. 16–23.
2. W.E. Burr, "Selecting the Advanced Encryption Standard," *IEEE Security & Privacy*, vol. 1, no. 2, pp. 43–52.
3. M. Clendenin, "WAPI Battle Exposes Technology Rifts with China," *EETimes*, 17 Mar. 2006; www.eetimes.com/news/semi/showArticle.jhtml?articleID=183700631.
4. W. Diffie and G. Ledin, translators, "SMS4 Encryption Algorithm for Wireless Networks," version 1.02, 8 May 2008.

Susan Landau *is a distinguished engineer at Sun Microsystems, where she works on security, cryptography, and public policy, including surveillance issues, digital rights management, and identity management. She is coauthor (with Whitfield Diffie) of* Privacy on the Line: the Politics of Wiretapping and Encryption, *updated and expanded edition (MIT Press, 2007). Landau has a PhD from MIT, an MS from Cornell University, and a BA from Princeton University. She is an AAAS fellow and an ACM distinguished engineer. Contact her at susan.landau@sun.com.*