

Trust and Legitimacy in Security Standardization – a new Management Issue?

Dirk Kuhlmann¹, Liqun Chen², and Chris J Mitchell³

^{1,2} Hewlett Packard Enterprise Laboratories, Bristol, UK

³ Royal Holloway, University of London, Egham, UK

Abstract. IT Security has to deal with a number of rather unique factors that pose new challenges for managing standardization processes in this field. This has recently led to attempts to establish this problem domain as an area of research in its own right under the heading *IT Security Standardization Research*. Research in this area focuses strongly on interdependencies between academic, institutional, and real-world practices, including aspects of governance. Such work represents one of several attempts to underpin trust in the processes, management, and results of IT security standardization, and may, in future, provide a hub for “externalized” and independent self-reflection in this area.

Keywords: IT security, standardization, vulnerability, exploit, transparency, review, trust, legitimacy

1. Introduction

Typical questions for standardization management involve the definition, control and streamlining of workflows that help ensure a maximum level of correctness, consistency, and topicality for standards. In contrast, questions regarding the trustworthiness and reliability of standards hardly appear on the radar. To a considerable degree this is due to how the institutional frameworks are set up, and who is trusted to enforce suitable editorial rules and access criteria for participation for work supporting these requirements. The inclusion of societal, political, or private concerns in technical standards is generally frowned upon, a few exceptions [1,2] confirming this rule. Confined to their purely technical role, typical standards tend to be treated as policy-neutral. This is not to say that their normative coverage and granularity, as well as their economic and legal connotations, is never a matter of serious contention. However, as far as the technical specification is concerned, the actual content of standards is typically considered to be beyond the impositions of politics, law, and public debate.

Why then might the legitimacy and trustworthiness, i.e. the public perception, of standards require dedicated management? This paper outlines why this kind of active intervention is necessary for IT security standardization. It has given rise to new fields of inquiry and to the establishment of dedicated research fora such as the “Crypto Forum” and the “Human Rights Protocol Consideration Research” working groups of the Internet Research Task Force (IRTF), as well as the Security Standardization Research (SSR) initiative presented in this paper.

2. IT Security as Matter of Fact and Matter of Concern

Mechanisms for preventing unauthorized access to, and safeguarding the correct functionality of, data processed on electronic devices have been a matter of major concern throughout the era of electronic data processing. As documented by the extensive coverage of IT-security related incidents in the general media and the ever growing security vulnerability databases, these concerns have become increasingly pronounced during the last decade.

In the not-too-distant past, the hunt for security flaws and their active exploitation primarily focused on the provider side, targeting servers or network infrastructure. Most vulnerabilities were treated as undesirable, and arose as unintended consequences of programming mistakes and sloppiness. This benign view of attributing security weaknesses to involuntary human errors has since changed, as has the choice of exploitable targets. Sophisticated attacks are now carried out against every type of device connected to a network, including embedded technology, industrial controllers, and end user devices.

Quite frequently, technical components marketed and deployed to provide protection introduce new security holes at the same time; some of these holes are even deliberately deployed to create an exploitable vulnerability. A seemingly infinite string of disclosures on the extent of digital surveillance increasingly fuels public concerns about the deliberate insertion of bugs and backdoors. Coding security mechanisms in a standards-conformant, robust and correct way is difficult, and it is common practice to reuse existing code. As a result, the impact of a vulnerability in a prominent implementation can be considerable. Weaknesses that affect specific products can rapidly produce a ripple effect if the vulnerability is in an underlying, more generally used, building block. Security protocols and encryption algorithms are prominent examples of such building blocks. Worse still are mistakes in the specifications and standards of IT security mechanisms, since these are likely to affect *all* existing implementations in equal measure.

It therefore comes as no surprise that IT security standardization has now come under scrutiny as a potential vehicle for undermining the security of systems and their unsuspecting users. And from here, it is not a large step to suspect the whole community of security specialists and standardizers as possible instruments of the ‘forces of darkness’. Such a catastrophic erosion of trust in the expert community is currently only a matter of concern, but active countermeasures might be necessary to prevent this concern from becoming a matter of fact.

3. Trust and Mistrust in IT security standards

What appears to set security standardization apart from many other fields is its sensitivity to errors: a small mistake can easily bring down the complete edifice. A given design is either secure or it isn't; this, somewhat peculiar, binary, nature of security means that an entire standard can go overboard in a blink of an eye. Problems are amplified by the existence of active adversaries. Deficiencies are not just exposed by the slow random walk of nature, but by well-organized endeavours aimed to unearth weaknesses. There can be few other areas of standardization where cohorts of individuals skilled in the art systematically probe the normative material for possible shortcomings. This effort can be compared to that of lawyers or accountants searching for legal or financial loopholes. This analogy carries over to the incentive structure: well developed markets exist in which security vulnerabilities are traded.

It is obvious that the success of IT Security standards relies on user trust. That is, potential adopters need to be confident that standardized schemes have been well-designed, and have not been deliberately manipulated to contain exploitable weaknesses. Similarly, standards writers need to be sure that contributions to the standards development process are well-founded. Both of these categories of trust have been seriously damaged by recent revelations about the deliberate inclusion of weakened cryptographic key management mechanisms in National Institute for Standardization and Technology (NIST) and ISO/IEC standards. Of course, the suspect mechanisms were de-standardised very rapidly, but the incident is continuing to have damaging effects on the development and use of standards, with all contributions from the US in particular being regarded as automatically suspect. This is hugely unfortunate, not least because over the last 40 years the US and NIST have played a major role in developing robust and useful security standards.

All this requires finding ways of moving beyond the current levels of distrust through ongoing dialogue, and to find better ways of gaining confidence in standards and proposals for inclusion in standards. Indeed, the whole issue of evaluating proposed standards is an area that needs much more work, and would benefit enormously from greater academic involvement, a key objective for the SSR conference series described further below.

4. Stakes and Stakeholders in IT Security

Another marked special characteristic of IT security standardization concerns its stakeholders. At least in our experience, IT standardization is typically influenced by rather a small circle of industrial players, frequently supplemented by a few interested academics. The large scale participation of members of specialist government departments is unlikely. While technological advances and market changes may sometimes shift the goalposts, an adjustment of goals rarely occurs as a reaction to external parties, i.e. self-appointed stakeholders not involved in the standardization process. The main success criterion for the typical IT standard is its incorporation in products that succeed in the market.

The playing field for IT security standardization is rather more chequered. Apart from the more typical promotion of the interests of individual stakeholders, and contributions from established academic circles, there is also the distinct possibility of undisclosed agendas for certain active contributors to the standardization process. Their interest might be to establish a security mechanism with well-hidden weaknesses, which prevent the masses breaking it, but enable the knowledgeable and well-equipped few. Again, problems are amplified by the fact that it is not blind nature that unveils mistakes; it cannot be assumed that flaws, once discovered, will be reported back to the standardizers. The lucky finder may instead choose to sell it on, or to retain it for future use (e.g. as a zero-day attack).

As security mechanisms become more diverse and complex, the number of individuals who fully understand them gets smaller. In IT security, it is not sufficient that specifications and implementations achieve the desired effects by following the normative description. Instead, a mechanism is only deemed fit for purpose when none of its side-effects can be exploited to undermine its original purpose. This property is much harder to demonstrate than mere functional correctness, as it requires a grasp of contextual parameters that must be given for the mechanism to perform correctly. Proper accounting for context sensitivity and comprehensive validation are underpinnings for the normative legitimacy of a IT security standard, which appears to mark yet another special characteristic.

The series of incidents referred to in sections 2 and 3 culminated in 2014 when a recommendation for generating elliptic-curve cryptography parameters, supplied by the National Security Agency (NSA) and endorsed by NIST, was withdrawn following years of critique. An investigation concluded that the selection criteria, security analysis, or any measures for quality assurance could not be reconstructed for the parameter generation process. There are indications that the provision of the questionable parameters for Dual_EC_DRBG was deliberate. A rushed process may have contributed to them being included in the standard [3].

The potential need to correct security standards to remove or patch vulnerable mechanisms, whether included by accident or design, raises a further issue, namely how to communicate the need for updates to all affected parties. It is impossible to know who has adopted a standard – permission does not need to be asked! As a result, methods need to be devised to disseminate the need for urgent changes as and when standards are revised.

The Dual_EC_DRBG incident has shown that even large, official standardization bodies may not command the capabilities to validate complex security mechanisms themselves, or may be unwilling to muster them. The research that was instrumental in building the case against the NSA/NIST recommendation was spearheaded by independent academics [4]. Close interactions between researchers and standardization bodies are likely to become ever more significant, both to find problems quickly and to help disseminate information about them. Improving these interactions has to be tackled as a management issue to be proactively addressed and driven by standardization practitioners and administrators. The Security Standardization Research (SSR) conferences presented below aim to provide an initial platform and framework to enhance these interactions.

5. Security Standardization Research in Context

By 2014, the need for an independent forum for IT security standardization stakeholders had been recognised by academic and industrial veterans, who had grown weary of an atmosphere often characterized by a level of mutual distrust and antagonism between corporate, governmental and academic participants.

The question was how best to address the lack of an adequate forum. Existing, well-established, IT security conferences are primarily interested in results of genuine novelty. Their modus operandi is very different to the drawn-out, consensus-oriented deliberations of IT security standardizers. Initiatives like the working groups of the IRTF Crypto Forum, on the other hand, require continuous participation. This is unattractive for researchers who may not only lack funding for such travel, but get little reputational mileage from the work. The IETF or IRTF is probably not a natural home for a new forum, since many relevant activities occur in official national and international bodies or industry consortia. Creating a conference series appeared to be the most promising avenue for bringing together the academic, industrial and institutional IT security camps. This setting encourages academic publication and invites contributions from IT security researchers who cannot participate in formal standardization efforts.

Establishing and maintaining a baseline of confidence and mutual trust between various stakeholders in IT security standardization has thereby become a matter for active management. A first step towards this goal has been to enable a more effective exchange of ideas between these groups. It took considerable effort, but was simplified by the fact that all the participants belong to the IT security expert community. However, improving interactions with the general public is a rather more difficult matter. Such interactions continue to be characterized by huge knowledge and information asymmetries, and it remains a fact that IT security continues to be unfavourably compared to traditional engineering disciplines. What is at stake here is the legitimacy of standardized IT security in terms of (a) designing and validating technical features and (b) responsible utilization. SSR was deliberately positioned not inside, but alongside, existing standardization institutions and processes, and there may be a chance to improve (a) through better transparency. So far, however, the SSR has made no determined attempts to address (b), i.e. the area of best practices and legal constraints.

6. First Practical Experiences

Most of the papers presented at the first two conferences fall into a small number of fairly well-defined categories, namely:

- evaluation, including formal analysis, of standardized security protocols and applications;
- evaluation of standardised cryptographic techniques;
- future topics for standardisation;
- privacy aspects of standardised protocols.

Both SSR conferences featured panel discussions with members from government organizations, industry and academia. The SSR 2014 panel was chaired by Joshua D Guttman (MITRE Corporation) and addressed “Formal Verification and Analysis of Protocols in Standards Development and Evolution”. In SSR 2015, Randall J. Easter (NIST) chaired a panel on “Accreditation, Validation and Recognition based on ISO Standards”. So far, at least two SSR papers have directly contributed to the improvement of standards, namely ISO/IEC 11770-4 [5] and ISO/IEC 11770 [6]. Representatives of major standards bodies were very happy to get involved in the SSR events, leading to very helpful and interesting discussions. The sessions revealed a need and willingness to discuss some of the key problems of security standardisation.

The conferences succeeded in getting academics more involved in discussions about the standardisation process, and those who came along left with a better appreciation of how and why standards are written. One objective of the SSR conference was to increase the academic involvement in security standards writing, and some success can be reported in this regard. An area that has so far been disappointing regards standards on security management. The ISO/IEC 27000 series standards have been very widely adopted; at the same time there has been quite a bit of vocal criticism of the ‘compliance approach’ to security. One hope was to get discussions of the advantages and disadvantages of the compliance approach, as well as possible alternatives, but so far almost all the contributions have been of a far more technical nature. There is clearly more work to do here, since the practical importance of the management standards is beyond doubt.

References

- [1] Internet Architecture Board. RFC 6973: Privacy Considerations for Internet Protocols. (July 2013). URL (28.02.2016): <https://tools.ietf.org/html/rfc6973>
- [2] Trusted Computing Group. TCG Design, Implementation, and Usage Principles (Best Practices) v3.0. (July 2011). URL (28.02.2016): http://www.trustedcomputinggroup.org/files/resource_files/5B50FA87-1A4B-B294-D0054DD2BACDF801/Best_Practices_Principles_Document_v3%200_Final.pdf
- [3] Jeffrey Carr. Six Cryptographers Whose Work on Dual EC DRBG Were Deemed Without Merit by RSA Chief Art Coviello. Digital Dao (Feb 26, 2014-02-26). URL (28.02.2016): <http://jeffreycarr.blogspot.dk/2014/02/six-cryptographers-whose-work-on-dual.html>
- [4] The corresponding Wikipedia page includes an extensive bibliography URL (retrieved 28.02.2016): https://en.wikipedia.org/wiki/Dual_EC_DRBG
- [5] F. Hao and S. F. Shahandashti. The SPEKE protocol revisited. SSR 2014.
- [6] C. Cremers and M. Horvat. Improving the ISO/IEC 11770 Standard for Key Management Techniques. SSR 2014.

Acknowledgements: We wish to express our gratitude and appreciation to all contributors to SSR-14 and SSR-15.