Security Compliance: The Next Frontier in Security Research

Klaus Julisch
IBM Research
Säumerstrasse 4
8803 Rüschlikon, Switzerland
kju@zurich.ibm.com

ABSTRACT
Practitioners as well as researchers have repeatedly deplored that IT security research has failed to produce practical solutions to growing security threats. This paper attributes this failure to the fact that IT departments no longer invest in security as an ideal. Rather, money is being spent on technologies that enable compliance with security requirements. Academia has not embraced this shift in perspective and still tries to "sell" security when organizations seek to "buy" compliance. This disconnect has lead to research that fails to improve real-world security because it is not embraced in the market place. The conclusion drawn in this paper is that academia needs to complement current security research by additional research into security compliance. To encourage more work in this relatively new direction, the paper describes the major compliance research challenges that await solutions.

Categories and Subject Descriptors
D.4.6 [**Operating Systems**]: Security and Protection; K.6.4 [**Management of Computing and Information Systems**]: System Management**.**

General Terms
Management, Economics, Security.

Keywords
Security, compliance, economics.

1. INTRODUCTION
It has repeatedly been deplored that security research failed to improve the security of society's IT systems [1, 2]. In the words of Bruce Schneier [3]:

> *"Computer security is a 40-year-old discipline; every year there's new research, new technologies, new products, even new laws. And every year things get worse".*

The reason for this failure lies at least partially in the way security is being practiced in commercial settings where the quest for security is increasingly being replaced by the quest for security compliance. *Security* is the state of being safe from threats (possibly subject to additional assumptions). By contrast, *security compliance* means conformance with a given set of security requirements, such as requirements on the use and configuration of virus scanners, firewalls, patching, penetration testing, or any other security mechanism. These requirements may be imposed by regulators, clients, or senior management and they determine what needs to be done to fulfill the security mandate.

While some may deplore this development, it is unlikely to revert because the reasons that brought security compliance to the center stage are unlikely to go away. Most importantly, these reasons include:

- **Audits and Enforcement:** What is being audited and enforced is compliance, not security[1]. Moreover, security incidents are tolerated more easily if one can show that they occurred despite the affected IT system being compliant with all applicable security regulations. Therefore, as long as careers are terminated and people go to jail [7] for failures in compliance – rather than security – the commercial world will continue to pursue compliance rather than security as their primary goal.
- **The Weakest-Link Phenomenon:** Security is only as good as its weakest link. A good security solution is therefore balanced and seeks to protect all systems components equitably and adequately. The implication of this is that practitioners are well-advised to seek compliance with some well-balanced security best-practice rather than chasing the latest security innovations, which generally strengthen a single link of the security chain. In other words, compliance makes sense from a practical point of view.
- **Measurement:** Despite extensive research [8,9,10,11], we are not capable of measuring security in the general case. Measuring compliance, by contrast, is feasible and several catalogs of compliance metrics have been published [12,13]. For example, "percentage of system with the latest patch level" is a compliance metric of the

NSPW'08, September 22–25, 2008, Lake Tahoe, California, USA.

[1] Even evaluation systems like TCSEC [4], ITSEC [5], or Common Criteria [6] evaluate the existence of specific security features and supporting assurance of their correctness. Such evaluations do not prove security or invulnerability to attacks and consequently, are more akin to audits that establish compliance with a set of security requirements. Section 2 offers a more detailed discussion of this point.

requirement to patch systems. Note, however, that this metric says little about the security of the system. It should not come as a surprise that C-level executives prefer a measurable discipline like compliance over an unquantifiable discipline like security.

These reasons have made *compliance with security requirements* rather than security itself the center of real-world security. As a consequence, organizations are in the market for "buying" compliance, and as long as academia "sells" security, the disconnect we observe today is likely to endure. Put bluntly, if the latest security widget out of our research labs is not required by any regulations or auditors, then nobody cares; if it improves the strongest link and leaves everything else unimproved, then nobody cares; and if said widget offers some unquantifiable improvement, then there is no "business case", and again, few commercial security professionals or their managers care.

This is not to minimize the importance of "classic" security research, but the implications from the above seems obvious: In order to narrow the gap between academia and practice, it is necessary to focus more research on the question of security compliance. Here, the central research question is: Given a set of security requirements (such as "deploy virus scanners A; use authentication mechanisms B; use monitoring and logging C; use patching policy D", and so on), how can one either build a new system or reconfigure an existing one so it provides those security controls as well as sufficient assurance of their correct functioning? Section 2 refines this definition of the security compliance problem, and Section 3 gives an overview of the interesting research questions it raises. Section 4 summarizes the paper.

2. THE SECURITY COMPLIANCE PROBLEM

The Webster Dictionary defines compliance as conformity in fulfilling official requirements. This paper obviously focuses on compliance with security requirements. IT security requirements can be classified into *functional requirements*, which require some functional security feature such as mandatory access control, and *assurance requirements*, which specify the evidence needed to establish that the functional requirements are met. With these amendments, I define:

Definition: Security compliance, in IT systems, is the state of conformance with externally imposed functional security requirements and of providing evidence (assurance) thereof. □

In a first approximation, the *Security Compliance Problem* then is the problem of becoming compliant with a given set of IT security requirements. The origin of these requirements can vary from government regulations to company-internal policies, but does not really matter for our discussion. Next, we will refine the definition of the security compliance problem.

As was indicated in the introduction, evaluation systems such as TCSEC [4], ITSEC [5], or Common Criteria [6] (which nowadays replaces TCSEC and ITSEC) test compliance with a set of security requirements. For example, the statement that a product has Common Criteria level EAL3 means that for a given set of functional security requirements, the product's compliance with those requirements has been established at confidence level 3, based on the available assurance. Common Criteria confidence levels are measured on a scale from 1 (lowest) to 7 (highest). It is worthwhile emphasizing that even though Common Criteria and the other evaluation systems do not use a "compliance language", they are really about compliance.

More recently, security compliance has been associated with regulations such as Title V of the Gramm-Leach-Bliley Act (GLB Act), the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), the Sarbanes-Oxley Act, the Data Protection Directive of the European Union, and others. The important difference between these regulations and the Common Criteria or other evaluation systems is that the regulations pose *new* security requirements on *existing systems*. By contrast, the objective to obtain Common Criteria certification for a product is generally known at design time, and the product is accordingly built to meet the desired EAL-level. In fact, for the EAL levels 5, 6, and 7 it is *mandatory* that security is built in starting at design time. Furthermore, Common Criteria have been used mostly for *products,* but evaluation of *systems* is less frequent and less well-defined. Making *existing systems* comply with *new security requirements* therefore is a new challenge of considerable practical importance. I therefore define:

Definition: Given an existing IT systems S and an externally imposed set R of security requirements. The **Security Compliance Problem** is to make system S comply with the security requirements R and to provide assurance that an independent auditor will accept as evidence of the compliance of system S with requirement R. □

The security compliance problem captures what "security" has become to mean in many real-world contexts and any researcher keen on having immediate real-world impact should be aware of this development.

The next section outlines the major research challenges that the security compliance problem raises. As will become apparent, these challenges are demanding and interesting and, in the author's view, worthy of more attention from the scientific community.

3. SECURITY COMPLIANCE FRAMEWORK & RESEARCH CHALLENGES

The general framework for solving the security compliance problem is a four-step process. Specifically, given an IT system S and a new security requirement R, the following steps have to be completed in order to make system S compliant with requirement R:

1. Formalize the requirement R.
2. Identify which sub-systems of S are affected by R.
3. Determine what assurance has to be provided to show that S is compliant with R.
4. Modify system S to become compliant with R and to provide the necessary assurance.

These steps and their associated research challenges are best illustrated by means of an example. For that purpose, let us consider an abbreviated version of Article 17 of the Data Protection Directive of the European Union [14]:

Article 17 [Abbreviated and simplified]: Member States shall implement appropriate technical and organizational measures to protect personal data against accidental, unlawful, or unauthorized destruction, alteration, or disclosure.

This requirement is typical of regulatory requirements and illustrates how vague such requirements tend to be. Without knowing what "appropriate technical and organizational measures" are and without a definition of "accidental, unlawful, or unauthorized" behavior, it is impossible to implement this directive. Step 1 of the above framework is intended to overcome this difficulty by eliminating the vagueness found in regulations. To date, this concretization and formalization process is mostly manual and while some researchers have started to address it [15], the need for better tools, languages, and methods still poses a largely unaddressed research challenge. An important aspect of this challenge is that our tools, languages, and methods must translate regulatory requirements into "good" security mechanisms, i.e. mechanisms *that actually improve security* rather than just fulfilling the spirit of some regulation.

Next let us assume it had been determined that "appropriate technical and organizational measures" means that a virus scanner has to be installed on every machine that processes personal data. This is clearly over-simplistic, but it is sufficient to illustrate step 2, which identifies the sub-systems that are affected by the security requirement. In our example, we must determine all machines that ever store or process personal data. This is a genuine challenge on many real-world IT systems, which tend to be complex and poorly documented. But localizing the data affected by a security requirement is only one of several challenges that step 2 has to overcome. Additionally, it needs to identify the affected software layers. In our example, just installing a virus scanner on machines is not enough. Rather, it is also necessary to configure the underlying operating system to protect the binaries and data files of the virus scanner or else the assurance is very weak.

In general, a security requirement can reach even the most unlikely sub-systems. Continuing with our example, it may, for instance, be necessary to modify boilerplate legal texts so they require business partners to install virus scanners on *their* machines. To the author's knowledge, there are no tools to support let alone automate step 2. This, clearly offers an opportunity for the research community.

Step 3 is concerned with determining assurance measures that are adequate or at least, sufficient to meet the auditors' demands. This step is still mostly manual and could benefit from automation. However, the probably larger research challenge is how to prove assurance in a trustworthy manner. This is the question of how an auditor can trust that the log files and measurements presented to him or her really come from the audited system and have not been modified. This question becomes even more important in dynamic Service Oriented Architectures and outsourcing environments where one service may want to check another service's assurance properties before engaging in a relationship. Research in this area has begun [16,17], but more work is needed.

Step 4 finally requires the necessary changes to be implemented in system S so it becomes compliant with requirement R. Given the complexity of today's IT systems, this is generally a demanding task. Several research questions arise: How would we have to engineer IT systems to make such a-posteriori changes, which are necessitated by new compliance requirements, easier to implement? What interfaces should software components implement to facilitate such changes? What compliance services – similar to security services [18] – can we implement in the IT infrastructure and middleware in order to standardize and simplify compliance functionality? Finally, are there policy languages that would allow us to formally express security requirements such as "install virus scanner on all machines processing personal data"? And if so, to what extend would such languages allow us to make compliance configurable?

4. SUMMARY & CONCLUSION

This paper has argued that from a business and practitioners' perspective, the quest for better security has become the quest for compliance with security requirements. The academic community has not yet embraced this shift in perspective and as a consequence, much of their research is not in demand among IT managers. Obviously, research cannot improve real-world security if it is not being used. To improve real-world security, we must therefore give the market what it wants, namely solutions to the security compliance problem. To stimulate more research in this area, the paper surveys the largely unresolved research challenges. This is not to minimize the importance of classic security research, but it is the author's contention that the disconnect between security research and practice can be narrowed if more researchers started to embrace security compliance as a research area.

Some may argue that it is a bad idea to focus more research on security compliance rather than security itself. After all, the countless hacker attacks, worms, viruses, stolen credit cards, compromised health records, and countless other breaches are very real and call for *very real improvements in actual security*. However, compliance with check lists of security requirements is *not* security and it is therefore only a distraction from the real goal of improving security. My response to this argument is three-fold: Firstly, this paper does not advocate to stop classic security research. Rather, I advocate to establish security compliance as a new research field. Secondly, helping organizations comply with a well-designed list of security best practices *does* – in fact – improve their security. Lastly, to improve real-world security, one has to be pragmatic and "sell" what organizations seek to "buy". For the foreseeable future, buying interest is for compliance solutions. As this paper has shown, there are many interesting research challenges in this area, which should help establish security compliance as a new research discipline.

5. ACKNOWLEDGEMENTS

The research leading to these results has received funding from the European Community's Seventh Framework Program (FP7/2007-2013) under grant agreement n° 216917.

6. REFERENCES

[1] E. Spafford, *Solving some of the Wrong Problems*, CERIAS Weblogs, 2007.

[2] N. Eppel, *Security Absurdity: The Complete, Unquestionable, And Total Failure of Information Security*, on http://www.securityabsurdity.com/, 2007.

[3] B. Schneier, *Testimony and Statement for the Record of Bruce Schneier*, on http://www.iwar.org.uk, 2001.

[4] Anonymous, *Trusted Computer System Evaluation Criteria*, Department of Defense , on http://csrc.nist.gov/publications/history/dod85.pdf, 1985.

[5] Anonymous, *Information Technology Security Evaluation Criteria (ITSEC)*, Department of Trade and Industry, London, 1991.

[6] Common Criteria, Part 1-3 on http://www.commoncriteriaportal.org/.

[7] W. Sturgeon, *Jail or compliance? You decide*, Directors Told, on http://www.silicon.com, 2005.

[8] B. Littlewood , S. Brocklehurst, N. Fenton, P. Mellor, S. Page, and D. Wright, *Towards Operational Measures of Computer Security*, Journal of Computer Security, 1993, pages 211 – 229.

[9] R. Ortalo, Y. Deswarte, and M. Kaaniche, *Experiments with Quantitative Evaluation Tools for Monitoring Operational Security*, IEEE Transactions on Software Engineering, 25(5), 1999, pages 633 – 650.

[10] E. Jonsson, *An Integrated Framework for Security and Dependability*, Proceedings of the 1998 Workshop on New security Paradigms, 1998, pages 22 – 29.

[11] M. Howard, J. Pincus, and J.M. Wright, *Measuring Relative Attack Surfaces*, Proceedings of Workshop on Advanced Developments in Software and Systems Security, 2003.

[12] D.S. Herrmann, *Complete Guide to Security and Privacy Metrics,* Auerbach Publications, 2007.

[13] NISP SP 800-55, *Security Metrics Guide for Information Technology Systems,* National Institute of Standards and Technology, 2003.

[14] The European Parliament, *Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data*, Official Journal of the European Communities, 1995 No L. 281 page 31.

[15] C. Giblin, A.Y. Liu, S. Mueller, B. Pfitzmann, and X Zhou, *Regulations Expressed As Logical Models (REALM)*, Legal Knowledge and Information Systems, 2005, pages 37– 48.

[16] M. Waidner, *Building Trust in Computing – Enterprise Security Perspective,* IST Conference, 2004, The Hague.

[17] J. Poritz, M. Schunter, E.V. Herreweghen, and M. Waidner, *Property Attestation – Scalable and Privacy-Friendly Security Assessment of Peer Computers,* IBM Research Report No. 99559, 2004.

[18] R. Kanneganti and P. Chodavarapu, *SOA Security*, Manning, 2008.