A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

> *"We need to establish security engineering as a valid profession in the minds of the public and policy makers. This is less about certifications and (heaven forbid) licensing, and more about perception – and cultivating a security mindset. Amateurs produce amateur security, which costs more in dollars, time, liberty, and dignity while giving us less – or even no – security."*
>
> Bruce Schneier
> Cryptographer and computer security specialist

Despite nearly 30 years of research and application, the practice of information system security engineering has not yet begun to exhibit the traits of a rigorous scientific discipline. As cyberadversaries have become more mature, sophisticated, and disciplined in their tradecraft, the science of security engineering has not kept pace. The evidence of the erosion of our digital security – upon which society is increasingly dependent – appears in the news almost daily.

In this article, we outline a research agenda designed to begin addressing this deficit and to move information system security engineering toward a mature engineering discipline. Our experience suggests that there are two key areas in which this movement should begin. First, a threat model that is actionable from the perspectives of risk management and security engineering should be developed. Second, a practical and relevant security-measurement framework should be developed to adequately inform security-engineering and risk-management processes. Advances in these areas will particularly benefit business/government risk assessors as well as security engineers performing security design work, leading to more accurate, meaningful, and quantitative risk analyses and more consistent and coherent security design decisions.

Threat modelling and security measurement are challenging activities to get right – especially when they need to be applied in a general context. However, these are decisive starting points because they constitute the foundation of a scientific security-engineering practice. Addressing these challenges will require stronger and more coherent integration between the sub-disciplines of risk assessment and security engineering, including new tools to facilitate that integration. More generally, changes will be required in the way security engineering is both taught and practiced to take into account the holistic approach necessary from a mature, scientific discipline.

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

Introduction

Despite nearly 30 years of research and application, the practice of information systems security engineering has not yet begun to exhibit the traits of a rigorous scientific discipline (Cybenko and Landwehr, 2012; tinyurl.com/kc3nm7p). As a result, it is still not possible to examine an information system and answer the question "How secure is it?" in a scientifically meaningful way. This is a significant problem because, increasingly, the economic and physical well-being of our society *depends* on the secure design and operation of business, government, and critical-infrastructure information systems. They appear in almost every facet of our daily lives but we actually know very little about how they stand up when it comes to security (Viega, 2012; tinyurl.com/mnwqd8c). It would be truly alarming to ask the question "How safe is it?" with respect to an aircraft only to discover that neither the engineers nor the certifiers really understood the answer. Yet, this is precisely the situation in which the information-technology security community finds itself today.

Although most of the concepts and ideas found in this article are applicable to security engineering at large, here we use term "security engineering" with specific reference to the security of information systems. Thus, in the context of this article, we define security engineering as "the art and science of discovering users' information protection needs and then designing and making information systems, with economy and elegance, so they can safely resist the forces to which they may be subjected" (National Security Agency, 2002; tinyurl.com/kcx5y4u). This definition has an analog in the physical sciences where the "forces" are natural and safety is an absence or avoidance of physical injury. As we shall see, this natural analog can inform us when answering questions such as "How secure is it?" in a more precise and consistent fashion.

Challenges

We see two significant challenges holding back the science of security engineering. First, it is unlike other engineering fields in the respect that the majority of the "forces" to be modelled are caused by human threat actors with *deliberate intent*, as opposed to forces due to natural and accidental causes. Thus, the first major hurdle facing security engineering is to define and maintain a threat model that can be used to calculate or bound these "forces" in a way that results in consistent engineering outcomes. This does not mean that threat models do not exist. In fact, like the nascent stages of any young scientific discipline, there are many models which, unfortunately, can lead to inconsistency and duplication of effort. For example, in many methodologies for assessing threats and risks, such as the Harmonized Threat and Risk Assessment Methodology developed by Communications Security Establishment Canada and the Royal Canadian Mounted Police (CSEC/RCMP, 2007; tinyurl.com/kfrjgv8) and the Guide for Conducting Risk Assessments developed by the National Institute for Standards and Technology (NIST, 2012; tinyurl.com/6srqlug), assessors are coached on *developing* a threat model. To be sure, threat analysts are not likely to rebuild their threat models each time they perform an assessment. Rather, the models are developed incrementally over time and are based substantially on individual knowledge and experience. However, while there may be commonality between models created by different assessors, there is certainly no guarantee that this is the case. This inconsistency (along with variations in categorization schemas, methodologies, definitions, and terminology) makes it challenging to validate and reuse results that would eventually drive the community to a small set of the most successful models. This convergence, which is a hallmark of a mature science, has not yet occurred within the security community.

Thus, we argue that a common threat model should be a primary goal of the security engineering community. This model should define the threat environment and the "forces" involved in a way that can be validated and built upon over time through repeatable qualitative or quantitative analyses. Such a model would also be "actionable" in the sense that threat-assessment results would point naturally to design options for security engineering that, at the outset, may be its primary measure of success. Such an undertaking would, of course, require a concerted research and development agenda to lay a common foundation upon which validation and refinement can begin to occur.

A second and potentially more challenging problem is the need for a useful framework for *security measurement*. Currently, there is no practical, relevant way of measuring the *absolute* security of an information system. In fact, there is no clear understanding of what absolute security means (e.g., Pfleeger, 2012: tinyurl.com/mt2xnsw; Böhme, 2010: tinyurl.com/kn99d4q; Davidson, 2009: tinyurl.com/l7475p3; Houmb et al., 2010: tinyurl.com/lw9ffqz; Savola, 2007: tinyurl.com/la87e84; Pfleeger, 2007: tinyurl.com/k3rjb6z; McHugh, 2002: tinyurl.com/m4z9nuu).

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

Absolute security is different from the common practice of "measuring" a system's compliance against arbitrary security requirements. Although a system may be 100% compliant with a set of security requirements, in most cases, there is little direct evidence that those requirements actually result in a more "secure" system. Clearly, security measurement will be assisted to a large extent by a common threat model, but the two approaches are co-dependent because threat modelling will eventually require quantitative measurement in order to demonstrate success.

We examine both of these challenges in the following sections and describe our vision for a security community-driven research program. For both challenges, we contend that the best approach is to take cues from established disciplines such as civil, mechanical, or electrical engineering and to draw analogies wherever possible. We feel that the closer we draw these parallels, the clearer will be our understanding of where we need to proceed next.

An Actionable Threat Model

A generally accepted model of deliberate threats is central to the advancement of the security engineering discipline. The most important aspect of such a model is that it be *actionable* from an engineering perspective. That is, when defining security requirements and undertaking risk-based design, the model would help to consistently and coherently identify a suite of design options – along with the associated security controls and their required level of implementation assurance – that could meet the goals identified by the system owner in terms of cost, operational utility, and risk tolerance. This is, of course, analogous to using standard engineering models during design activities (e.g., high-frequency antenna design).

Typical methodologies for assessing threats and risks, such as the Harmonized Threat and Risk Assessment Methodology (CSEC/RCMP, 2007; tinyurl.com/kfrjgv8) and the NIST's Guide for Conducting Risk Assessments (2012; tinyurl.com/6srqlug), generally focus on the generation of information related to risk decisions. In these assessments, a potentially long list of threat actors or threat scenarios is generated to estimate threat attributes and calculate the potential for risk. In other words, motivation is assessed for the purposes of determining the likelihood of an attack. Few, if any, threat attributes are identified and assessed in a way that purposely helps with the selection of design options or security controls, or that helps determine levels of implementation assurance.

In order to achieve an actionable threat model, there are two fundamental changes that must be made to the way threats are assessed. First, the act of performing a *threat assessment* must be divorced from the act of performing a *risk assessment*. Existing threat and risk-assessment frameworks make little distinction between these activities. Although they have elements in common, each activity requires a different skill set and targets different audiences. Second, threats should be assessed based (at least initially) on the *capabilities* that a threat actor *could* wield rather than on attributes that are specific to threat actors themselves (e.g., motivation, intent, risk aversion, willingness to invest time). Actor-specific attributes are more appropriately addressed during risk assessment. We explore a capabilities-based approach in the following sections.

Threat assessment based on threat actor attributes

Typically, threat-assessment methodologies begin by asking which threat actors are likely to attack an information system. In some cases, threat actors may be dropped from consideration if the likelihood of an attack is deemed to be very remote. More often, this likelihood is used to condition the potential risk from the attacker downwards (the injury from an attack does not change, just the magnitude of the outstanding risk). This approach places bounds on the costs of security, both in terms of money and constraints on operational freedom, and focuses limited resources where significant injury is *expected* to occur.

The likelihood that a threat actor will attack is often determined by examining certain attributes such as the actor's capabilities (what kinds of attacks they are capable of doing), motivations and intents, aversion to risk, willingness to invest time and effort, degree of access, etc. A difficulty with this approach is that many of these attributes cannot be accurately modelled or assessed because they can change frequently over time or they are based on complex mental states or behavioural patterns. As a consequence, assessments based on these attributes have large uncertainties, which makes the expectation of where significant injury will occur less accurate.

A more significant challenge with this approach is that an analyst must develop an exhaustive list of credible threat actors and their attributes in order to ensure that all threat scenarios are addressed. However, it is difficult to reason about the completeness of this list as

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

demonstrated by such events as the Oklahoma City bombing (tinyurl.com/gesg2), the September 11 attacks (tinyurl.com/nhx7m), the Fukushima Daiichi nuclear disaster (tinyurl.com/44hjgf3), and the Lac Mégantic derailment (tinyurl.com/q2fh9nk).

In the course of compiling this list of credible threat actors, the analyst must also think about, enumerate, and assess the threat actor's *capabilities* – the ways in which each threat actor might attack the system. As illustrated in Figure 1, the challenge with this approach is that capabilities can be missed if, for example, the assessor fails to consider every possible threat actor scenario or if threat actors are not adequately considered because they are viewed as being unlikely to attack. Some of these unidentified capabilities, if exercised, could be crippling to an organization. This revelation will not necessarily be made obvious by thinking only about threat actors. It is also natural to assume that even those threat actors that *have* been considered will evolve over time and that some may come into possession of more sophisticated, or even as-yet unidentified capabilities. Only a well-disciplined, frequent refresh of the assessment of threat actors will be able to track this evolution. In the next section, we argue that a better approach is to base a threat assessment on capabilities instead of threat actors themselves.

Threat assessment based on threat actor capabilities

Instead of modelling the characteristics of threat actors such as their motivation and intent, resources, and tolerance to risk, we propose that the threat assessment should be focused on the *capabilities* that can be employed to attack a system. Using this approach, it is possible (although potentially challenging) to: i) develop a more exhaustive survey of the threat *potential*, ii) reason about the completeness of the analysis, and iii) identify potential gaps in our knowledge. A capabilities-based approach also lends itself to a community effort because system-specific information need not be divulged. Given that capabilities can be assessed in a general context, the material will be highly reusable. A final benefit of this approach is that it provides a common interface between attacks by threat actors and the controls necessary to effectively counter them, as illustrated in Figure 2.

But what do we mean by threat actor capability? In our view, a *capability* is composed of a *vector* and a *sophistication level*. In simple terms, a *capability vector* defines a coarse taxonomy of attacks on an information system (Figure 3). A capability vector identifies *where*, *how*, and *what*. These fields of a capability vector can include:

1. **Access Mode:** This field identifies the means by which access to the target system is obtained. *Direct modes* include physical, personnel, logical, and electromagnetic access. *Indirect modes* involve direct modes that are applied to lifecycle elements of the system (e.g., system development, software patches, replacement hardware, and support-system operations). Indirect modes are potentially recursive and generally relate to the supply chain of a system.

2. **Target Layer:** A capability vector will be designed to act against one or more architectural "layers" within a system, such as the application, data, operating system, virtualization, network, firmware, and hardware.

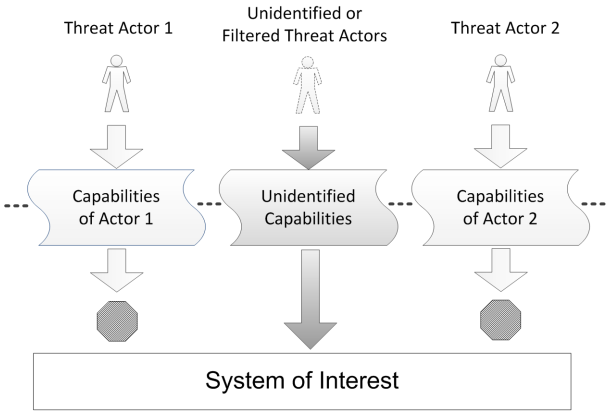

Figure 1. Capabilities from unidentified threat actors may be overlooked

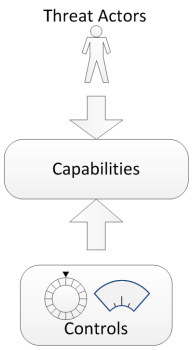

Figure 2. Threat actors are related to controls through their capabilities

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

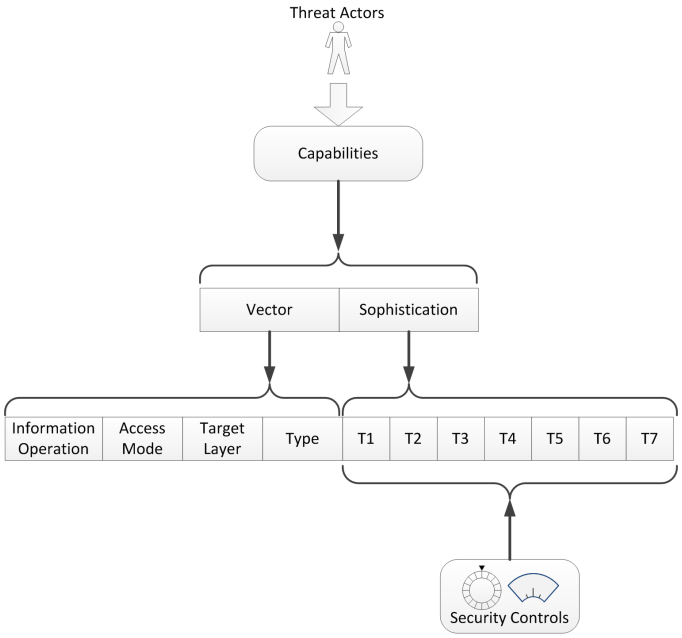

Figure 3. The threat capability model

3. **Type and Sub-type:** These fields identify a common name or grouping for a specific capability in order to maintain semantic compatibility with common terminology for various attacks. The Type field may be followed by a number of Sub-type fields to further distinguish capabilities. An example of the material that could be found in the Type field could include top-level categories found in the Common Attack Pattern Enumeration and Classification (CAPEC; capec.mitre.org).

4. **Information Operation:** This field is optional but can be helpful as a way of grouping or indexing attack vectors that have similar effects. The Information Operation field hints at *why* a given capability might be exercised (i.e., its intent). A capability vector may belong to more than one type of information operation. Potential categories include deny, exploit (infiltrate, exfiltrate), reconnoiter, deceive, etc.

It should be noted that this breakdown is only a recommended starting point. More, fewer, or different categories may be needed as the framework evolves.

The second element of a capability is defined by its *sophistication*. For example, consider a denial-of-service capability. A distributed version of this capability can be performed using software downloaded from the Internet and executed from a dozen computers. The same capability can be launched from 5000 computers distributed all over the world using code that exploits a previously unknown vulnerability. The differences between these capabilities are: i) the level of sophistication required to set up and execute them and ii) the set of controls required to prevent or limit the successful use of the capabilities against a system (as well as the rigour with which they are designed, implemented, and operated).

Thus, simply identifying a capability vector is not enough. To fill out the threat assessment, security assessors must also determine if there are distinguishing features (or attributes) that make the same capability vector harder to detect or prevent and then identify what options exist to address these more sophisticated variants (i.e., controls, architecture changes). In Figure 3, we divide sophistication into seven distinct levels according to those originally proposed in the National Security Agency's Information Assurance Technical Framework (NSA, 2002; tinyurl.com/kcx5y4u). However, we have not yet determined what would constitute a worthwhile set of distinguishing sophistication attributes, although we suspect that they may be somewhat dependent on features of each individual capability vector.

It is important to emphasize that identifying graduated levels of sophistication leads to the selection of security controls and design options that are generally more expensive or more operationally constraining as one moves up the scale of sophistication. This approach provides risk assessors with more explicit information regarding the tradeoffs between threat mitigation and costs.

As a final note, an important feature of the capabilities-based approach is that it has some predictive potential. That is, if we expanded every combination of the first attribute with the second and then third attribute, we obtain the universe of possible capability combinations. Some combinations will not make sense and can be disregarded while others will have ample evidence to show that they are in active use by threat actors at various levels of sophistication (e.g., logical access, operating system, Trojan, infiltration). Other combinations will appear strangely unfamiliar, either because they have never been exercised or they have been exercised but have never been publicly observed (e.g., electromagnetic, hardware, audio covert channel, exfiltration). Thus, capability vectors should tell us where we need to be looking for evidence of attack and, as a corollary, where threat actors might look in order to find new opportunities to expand their capabilities.

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

Figure 4 illustrates a fictitious set of capabilities related to a denial operation. For each capability vector, there are seven cells in which information about sophistication can be captured. Exactly what information should be contained in each cell and how it is determined is a fundamental research problem. To satisfy risk assessors' need to quantify the "potential" for an attack using a given capability, we propose three general categories as follows:

1. The capability has been observed to be in use by at least one threat actor at the given level of sophistication (i.e., the black boxes marked with an "O" in Figure 4).

2. The capability has been demonstrated at a conference such as DEF CON (defcon.org), but has not yet been observed "in the wild" (i.e., the hashed grey boxes marked with a "D" in Figure 4).

3. The capability is known to exist at a given level of sophistication but has not been observed (i.e., the dark grey boxes marked with an "E" in Figure 4); an example would be a nuclear-generated electromagnetic pulse.

Many capabilities follow a "commoditization" lifecycle in which they are generated at high levels of sophistication but are subsequently made easier to implement and become more widely available at lower levels of sophistication. This can be represented in Figure 4 as a heat map and could provide valuable information for risk assessors when considering the need for security controls over a long period of time.

From an engineering perspective, each cell should map to a set of security controls, mechanisms, strategies, security design patterns, and implementation-assurance requirements that have been shown (preferably through quantitative analysis) to effectively counter the threat capability at the specified level of sophistication.

Regardless of the type of information included with each cell, the basic research problem is as follows: given a post-incident analysis of a threat event (or an analysis based on vulnerability research work), how do we determine what sophistication level is represented? We see this as a difficult challenge because it will be a necessarily subjective exercise (at least at the outset). However, there are both theoretical and practical approaches that can help to reduce variation and uncertainty introduced by this subjectivity.

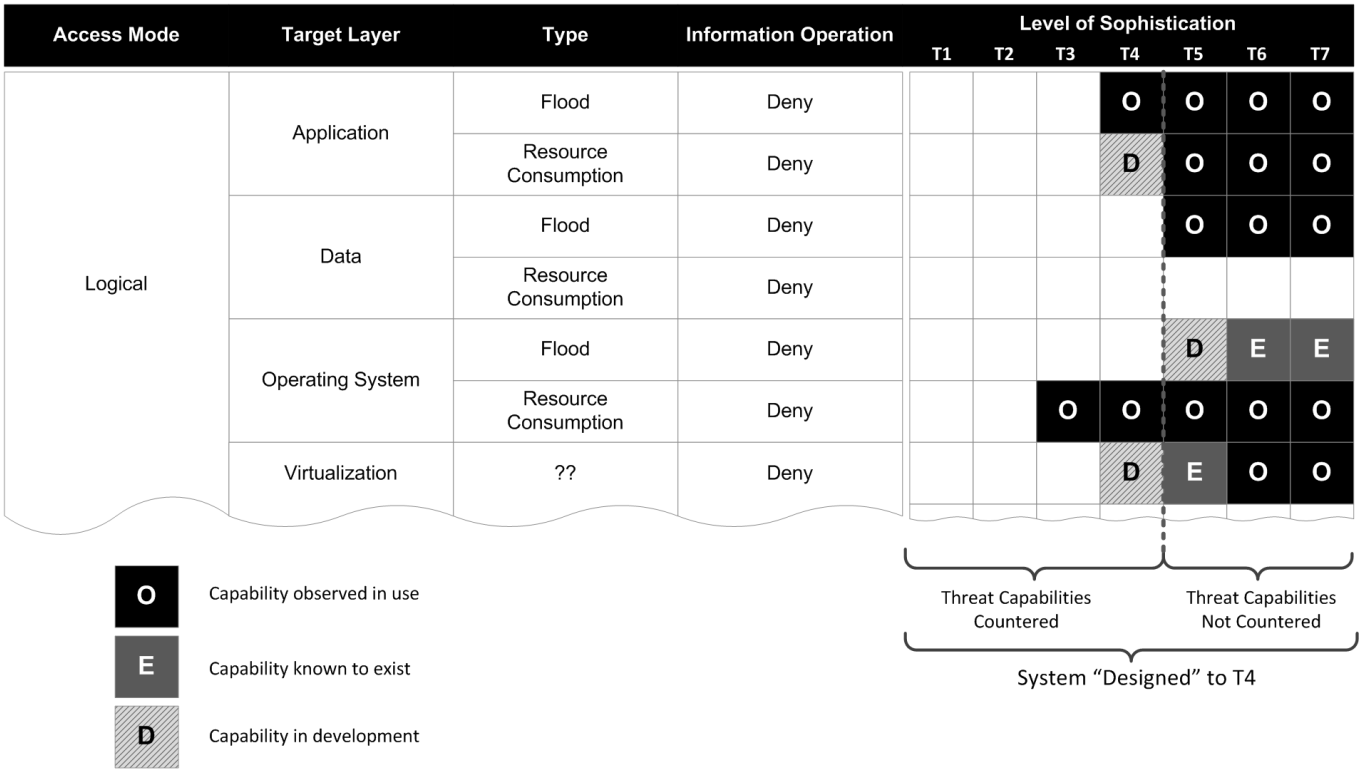

Figure 4. A threat-capability example for a fictitious set of vectors

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

Security Measurement

A second key challenge facing information-system security engineers is the difficulty involved in actually *measuring* security in a *practical* and *relevant* way (Pfleeger, 2007; tinyurl.com/k3rjb6z). There is a significant body of research on the subject but all studies seem to fail one or both tests of practicality or relevance (e.g., Lundin et al., 2006: tinyurl.com/mez2k8k; Shin et al., 2011: tinyurl.com/k2nvk53). If this were not the case, we would have a working solution by now. Being able to measure security in a useful way is absolutely critical to the advancement of security engineering as a discipline, because measurement is the bedrock of the scientific approach.

Why measurement is difficult

Security measurement is challenging for a number of reasons. First and foremost is the problem of concisely defining what is meant by the term *security*. Krautsevich, Martinelli, and Yautsiukhin (2010; tinyurl.com/kpuek8j) note that "we do not have a widely-accepted and unambiguous definition" that enables us to indentify one system as more secure than another. However, the definition of security often depends on perspective and context; it means different things to different people in different roles. Thus, we believe that there are a multitude of definitions that may all be equally useful within their own contexts. For example, contradicting security objectives may arise when considering an organization's need to monitor and control what happens on their systems versus an employee's need for legal privacy protection. Nevertheless, because the definition of security will have a significant impact on the way it is measured, it is critical to ensure that it is chosen appropriately and used consistently.

A second difficulty is that, regardless of how they are defined, security properties must actually be *measurable* and those measurements must be *practical* to obtain. There are at least three types of *security measurement* that information-system security should be concerned with:

1. **Engineering measurement:** These measurements are used by engineers to build models that "provide a formal representation (e.g., sets of equations) that corresponds well to security for systems under consideration" (Verendel, 2010; tinyurl.com/lgsxnrl). These are the same kind of measurements that one would expect from, for example, stress and strain analysis of various materials in civil engineering.
2. **Compliance measurement:** These measurements establish the degree to which an information system meets a set of specifications derived from security functionality and assurance requirements. Compliance measurement is normally performed throughout the process of system development. Examples of these types of measurements are described in the "Overview of IT Security Risk Management: A Lifecycle Approach" (CSEC ITSG-33: Annex 2, 2012; tinyurl.com/kf5ejyu) and the Common Criteria (commoncriteriaportal.org).
3. **Operational measurement:** These types of measurement provide metrics to reflect the operational security performance of an information system. Examples include patch-management coverage, mean time to mitigate, etc. Related resources include the ISO/IEC 27004:2009 standard for measurement techniques in information security (tinyurl.com/ln92xe3) and the metrics used by the Center for Internet Security (cisecurity.org).

In this article, we are concerned primarily with engineering measurement because it is a prerequisite for advancing the science of security engineering. Unfortunately, these kinds of measurement also appear to be the most difficult to obtain in a *quantitative way* (Wang, 2005: tinyurl.com/mgj3mj3; Verendel, 2010: tinyurl.com/lgsxnrl). They require a common, objective scale and a measuring device, and both must be accepted broadly across the security-engineering community in order to gain traction (Zalewski et al., 2011; tinyurl.com/lqw6865). Security engineering lacks these quantitative standards, primarily because security is often expressed in abstract terms.

In the absence of quantitative measurements, qualitative assessments have been used to derive security metrics. Qualitative assessments may be the best that security engineering can achieve until appropriate quantitative measures become available. Unfortunately, subjectivity implies inconsistency, which is unacceptable in a science-based discipline. Although it may not be possible to eradicate subjectivity altogether, there are certainly ways to minimize it. In some respects, we are advocating the same approach (but on a much larger scale) that the National Institute of Standards and Technology (NIST) has taken with cryptography. NIST arranges encryption algorithms by key size according to number of "bits of security" that they provide. The scale is *nominally* objective but an algorithm's placement on the scale is the result of expert judgment by one or more

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

cryptographic mathematicians who estimate the "amount of work" necessary in order to crack a ciphertext using an algorithm with the given key length (NIST 800-57: Revision 3, 2012; tinyurl.com/n5dk85u). Here, the measurement is a subjective assignment on a constructed scale and clients use that as an engineering measurement to compose their designs.

A third challenge facing security measurement is the notion of *assurance*. In the cryptography example above, we looked at measurement from the perspective of answering "How strong is it?" but we have not asked the question "How well does it work?" Without addressing assurance, efforts in addressing security are wasted. For example, although an encryption algorithm may have strong *conceptual* security (i.e., theoretical strength), if the algorithm is implemented incorrectly, then its *actual* security (i.e., robustness) is weak.

Assurance measurement is not as widely considered as *strength* within the security research community. Notable exceptions include cryptographic evaluations following the Federal Information Processing Standard (FIPS 140; tinyurl.com/pn9mb4) and Common Criteria assurance requirements (commoncriteriaportal.org/cc/). This lack of attention to assurance measurement is probably due to the fact that it appears to be even more abstract and (in most cases) more subjective than measurements of security strength. Generating and communicating assurance information in a "standardized" way would serve to reduce subjectivity, and this is the focus of at least one object-management group's specification (Alexander et al., 2011; tinyurl.com/mtjuufn). However, combining information about assurance and strength into a composite measure of security should be subjected to further analysis and validation. Some work has been accomplished in this direction with the notion of "robustness" in the CSEC's IT Security Guidance (CSEC ITSG-33: Annex 2, 2012; tinyurl.com/kf5ejyu), the Information Assurance Technical Framework (National Security Agency, 2002; tinyurl.com/kcx5y4u) and some Common Criteria protection profiles (commoncriteriaportal.org/pps/).

Going forward

On the issue of pursuing a research agenda that addresses practical and relevant security measurement, we propose a straightforward approach: consider security engineering in relation to other mature engineering disciplines and draw as many analogies as possible. This approach has been advocated before by Zalewski and colleagues (2011; tinyurl.com/lqw6865), who perceive that moving ahead requires "a closer alignment of security assessment with concepts developed in measurement science and physics". Where an analog in security engineering does not exist or does not translate easily, then we have an item to add to the research agenda. We suspect this might help identify the form that measurements *should* take.

For example, civil engineers build structures that are designed to exist in a certain threat context. "Loads" (in terms of forces) are applied in three-dimensional space – often downward with the pull of gravity but sometimes in other directions due to other natural or man-made forces. The scale and magnitude of these loads is directly proportional to defined levels within each threat event; for example, an "F2 tornado" (tinyurl.com/2frdj2) or a "Cat3 hurricane" (tinyurl.com/kl5ukgo). Minimal load levels for certain threat events are specified by regulatory bodies and have an effect on the way the structure is architected and designed (e.g., minimal spacing between load bearing members). Other, unregulated threat events may be of specific concern to certain clients and the forces to be countered by these threats may be specified as additional design requirements (e.g., bollards in front of federal buildings).

A natural security analog to "loads" is provided by the spectrum of threat sophistication that we proposed earlier. However, although we imply that the "load" imposed by a threat capability vector at sophistication level 7 is "greater" than level 6, we do not have a clear understanding of what "forces" are applied by threat agents against the information-system infrastructure and what effects these may have. Addressing these gaps in our understanding may help us develop the engineering metrics that are needed to advance the science of security engineering.

Putting It All Together

In the following sections, we outline a few specific areas where improvement can be expected as a result of taking on the research agenda proposed in this article.

Composite security

The "holy grail" of security engineering is to be able to answer the *composition problem* (Irvine and Rao, 2011: tinyurl.com/n626rgo; Datta et al., 2011: tinyurl.com/kukg98y). That is, given an information-system architecture or design made up of discrete security and non-security components, solving the composition problem would allow us to determine the overall security of the information system. The composition problem is a common lament in the information-security domain; "We simply have no theoretical basis for judging the security of a

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

system as a whole" (McHugh, 2002; tinyurl.com/m4z9nuu). However, the composition problem cannot be solved without measurement, and measurement cannot be performed without a generally accepted threat model.

Development of mandatory security requirements
Having a science-based threat model and security-measurement framework would allow the security community to influence the development of security standards that are based on sound engineering principles. In civil engineering (and certainly other engineering disciplines), threat events that may pose a risk to safety are incorporated over time into standards, codes, and regulations. This information is gleaned from engineering measurement and, in some cases, spectacular failures such as the Tacoma Narrows Bridge (tinyurl.com/c77rpw). A hallmark of a mature engineering science is the ability to investigate and learn from these failures and recycle that information into curricula, codes, and regulations.

Security engineering appears to have few close equivalents to requirements specified in codes and regulations – anti-virus and access control mechanisms seem to be a standard requirement found in most system specifications, although these are by no means mandated. In order to begin embedding security controls in security standards (especially if they are very expensive), it is necessary to thoroughly understand those controls from an engineering-science perspective.

Security-engineering curriculum
Finally, we note the fact that many curricula being proposed for security engineering in a college or university setting are simply computer engineering or computer science degrees that have been sprinkled with topics in security, assurance, and, unfortunately, risk assessment or risk management (e.g., Hjelmås and Wolthusen, 2006: tinyurl.com/kncwfek; Older and Chin, 2012: tinyurl.com/l5bbtah; Irvine and Nguyen, 2010: tinyurl.com/mvzj4xa). As far as we know, there is no curriculum that seeks to build (or build upon) a set of mathematical (or at least more formal) models that allow the composite security of an information system to be determined in a repeatable, meaningful manner. We suspect this is due to a lack of understanding of where exactly to begin.

Conclusion

In this article, we broadly outlined a research agenda that, with sufficient effort, would help begin the process of placing security engineering for information systems on foundations equivalent to other mature engineering disciplines. Two significant areas requiring attention were identified: threat modelling and engineering-security measurement. We argued that these areas are critical starting points because they affect almost all other aspects of security engineering, and more generally, the field of IT security. In addition, we believe that in order to be successful, these areas of research should be performed by a multi-disciplinary team of subject-matter experts. In taking on this research agenda, there is considerable opportunity to affect a significant change in the security posture of existing and future information systems. And, in doing so, the security and privacy of Canadians and the trust that they invest in the information systems of businesses, governments, and critical-infrastructure information systems will also be positively affected.

Acknowledgements

We would like to thank Larry Stoddard for his contribution and insights into this work.

A Research Agenda for Security Engineering

Rich Goyette, Yan Robichaud, and François Marinier

About the Authors

Richard Goyette is Senior Security Architect at Communications Security Establishment Canada. Richard has a BEng and MEng in Electrical Engineering, both from the Royal Military College of Canada in Kingston, Canada. Richard spent 22 years as a Signals officer in the Canadian Forces, where he was involved with a multitude of projects in the areas of intelligence, security, and command and control. He is currently employed in the area of architecture and technology assurance developing security guidance for the wider Government of Canada.

Yan Robichaud is a Senior Security Architect at Communications Security Establishment Canada. Yan has a BASc degree in Computer Engineering and MSc degree in Electrical Engineering, both from Université Laval, Québec City, Canada. He provides advice and guidance related to security architecture and engineering, threat assessment, and risk management to Government of Canada departments and agencies. He is involved in key government IT initiatives, such as large IT consolidation projects, enterprise security architecture, and the security of space-based systems. Yan is also involved in the development of IT security courses and leads the production of publications about IT-security guidance, such as "ITSG-33 IT Security Risk Management: A Lifecycle Approach".

François Marinier is an independent IT security analyst with experience in all facets of IT-security risk management. François started his career working in computer operations and mainframe application support. He eventually migrated to IT security, where he acquired knowledge and experience in the development and application of processes for IT-security risk management. He has also worked as an analyst, supporting large IT-infrastructure initiatives, in both the public and private sectors. For the last three years, François has dedicated his work almost exclusively to the development of ITSG-33, the next generation of guidelines for IT security risk management for the Government of Canada.

Citation: Goyette, R., Y. Robichaud, and F. Marinier. 2013. A Research Agenda for Security Engineering. *Technology Innovation Management Review*. August 2013: 41–50.

Keywords: cybersecurity, information system security engineering, research, risk management, security engineering, threat modelling, security measurement