Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

> *"Risk comes from not knowing what you're doing."*
>
> Warren Buffett
> Business magnate, investor, and philanthropist

Progress in operational cybersecurity has been difficult to demonstrate. In spite of the considerable research and development investments made for more than 30 years, many government, industrial, financial, and consumer information systems continue to be successfully attacked and exploited on a routine basis. One of the main reasons that progress has been so meagre is that most technical cybersecurity solutions that have been proposed to-date have been point solutions that fail to address operational tradeoffs, implementation costs, and consequent adversary adaptations across the full spectrum of vulnerabilities. Furthermore, sound prescriptive security principles previously established, such as the Orange Book, have been difficult to apply given current system complexity and acquisition approaches. To address these issues, the authors have developed threat-based descriptive methodologies to more completely identify system vulnerabilities, to quantify the effectiveness of possible protections against those vulnerabilities, and to evaluate operational consequences and tradeoffs of possible protections.

This article begins with a discussion of the tradeoffs among seemingly different system security properties such as confidentiality, integrity, and availability. We develop a quantitative framework for understanding these tradeoffs and the issues that arise when those security properties are all in play within an organization. Once security goals and candidate protections are identified, risk/benefit assessments can be performed using a novel multidisciplinary approach, called "QuERIES." The article ends with a threat-driven quantitative methodology, called "The Three Tenets", for identifying vulnerabilities and countermeasures in networked cyber-physical systems. The goal of this article is to offer operational guidance, based on the techniques presented here, for informed decision making about cyber-physical system security.

Introduction

Cyberattacks are increasing in frequency and severity. Prolexic Technologies (2013; tinyurl.com/n66algm) reports that the average packet-per-second rate of distributed denial-of-service (DDOS) attacks reached 47.4 million packets per second and the corresponding average bandwidth reached 49.24 Gbps in the second quarter of 2013. These are increases of 1,655% and 925% respectively over 2012.

Although DDOS attacks are relatively brutish cyberweapons, the so-called "advanced persistent threat" (APT) refers to sophisticated attackers who operate more subtly against specific targets with specific goals. For example, Operation Aurora deployed a zero-day web-browser exploit to extract detailed intellectual property from high-tech companies (McAfee Inc, 2010; tinyurl.com/np89339).

Whether done with blunt objects (DDOS) or scalpels (APT), cyberattacks continue to be effective. In fact, enterprise IT security managers believe their networks are becoming less secure. A survey of 671 IT security practitioners conducted by the Ponemon Institute (2012; tinyurl.com/afk94px) found that only 33% believed their IT

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

networks were more secure in 2012 than in 2011. In spite of such concerns, a recent Oracle study (2013; tinyurl.com/l76858h) found that, even with increased overall IT security spending, enterprises are still not protecting the right assets.

Combining all these facts and findings, it is evident that the growing cyberthreat environment is becoming more complex and more targeted while our ability to respond with appropriate defences at the appropriate investment levels is becoming more difficult.

The cybersecurity research, development, and vendor communities have not been helping matters. Most researchers and vendors promote their specific point solutions at the expense of seeing the bigger security picture. For example, on the one hand, the "build security in" community advocates redesigning and rebuilding IT systems from scratch to be more secure from the start (e.g., U.S. Department of Homeland Security: tinyurl.com/mh4a2e3; Darpa: tinyurl.com/6nf5yp3; McGraw, 2013: tinyurl.com/mu4oz24). On the other hand, "big data" security technologies promote extensive IT instrumentation, logging, and analysis for whatever application and network infrastructure that has already been deployed (e.g., Hewlett Packard: tinyurl.com/kdsrvuj; Splunk: tinyurl.com/kj3ujkp).

These extremes beg the key question of what combination of cyberdefensives are appropriate for securing an enterprise from the spectrum of threats that it realistically faces. Government efforts at articulating security best practices and risk assessments (e.g., National Institute of Standards and Technology, 2013: tinyurl.com/c8vukj7) are comprehensive and noble but too generic to be operationally prescriptive for such purposes.

New ideas are needed for enterprise-level cybersecurity assessment and investment. The novel approach proposed in this article is based on the authors' 30 years of combined experience in securing complex cyber-physical systems in government and private sector environments. The approach consists of three ingredients that will be outlined in detail below:

1. Confidentiality, integrity, and availability requirements
2. Quantification and assessment of cybersecurity defence investments
3. Identification of cybersecurity threats and vulnerabilities

This article is organized around these ingredients, as follows. The second section argues that tradeoffs between confidentiality, integrity, and availability are intrinsically unavoidable in typical enterprise operations and proposes an analytic framework for managing those tradeoffs. The third section describes a methodology for quantifying the impact of vulnerabilities and defences that are used to mitigate them, namely "QuERIES". The fourth section presents the underlying cybersecurity model, called "The Three Tenets" of cybersecurity-vulnerability assessment and mitigation. Finally, the fifth section provides a summary and a discussion of ways forward based on these results.

Confidentiality, Integrity, and Availability Requirements

Security considerations and metrics are not the only criteria enterprise IT managers use to make decisions. Revenue (or service in the case of a non-profit or government entity) is a result of providing users access to networked services and information and so it is often a primary driver when trading off security against access.

In practice, decision makers must constantly balance availability (i.e., the ability of end users to derive benefit from the system), confidentiality (i.e., the protection of information from access by unauthorized users), and integrity (i.e., the protection of information from unauthorized modification). This task involves complex, typically enterprise- and system-specific, tradeoffs that require an appropriate balance between properties that are not entirely consistent with each other.

In order to make such tradeoff decisions more rigorous and quantitative, we have started to develop a model and corresponding framework for confidentiality, integrity, and availability (CIA) risk management. Here, we briefly introduce our work on the specific issue of introducing "diversity" into an enterprise IT environment for the purpose of increasing "security". Information-system diversity, as opposed to "monoculture", has often been praised as a mechanism for building more resilient and secure systems, ones in which the compromise of one system does not immediately translate into the subsequent compromise of all similar systems.

Diversity can be introduced into an IT system by deploying hardware and software from different vendors or by mechanisms such as randomizing address layouts or compiler generation of executable code (Jajodia et al., 2011; tinyurl.com/mz5d8fn). Further details of the model and associated results can be found in a forthcoming

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

technical paper devoted to this issue (Cybenko and Hughes, *in press*; tinyurl.com/m3jexfv).

Our basic model assumes a network of nodes that comprise an asynchronous distributed system that an enterprise operates. These nodes could be mirrored web or database servers, clients, routers, or other replicated devices or services in an information system. The designer has a choice of making the components the same (i.e., homogeneous or a monoculture) or making the components different in some way (i.e., diverse, moving targets, heterogeneous, or some other approach).

A compromise of a node (component or device) means that an attacker has control of that node, such as root or administrator privileges in an operating system, for example. The goals of a compromise are often summarized as violating one or more of the confidentiality, integrity, and availability properties (Smith and Marchesini, 2008; tinyurl.com/l8jx7op). We interpret these goals in the context of a networked system of functionally redundant components. In this article, we define these concepts as follows:

- *Availability* means that at least one of the nodes has not been compromised and is therefore functioning properly. Stated otherwise, not all of the nodes in the system have been compromised and so at least one is still functioning in a reliability theory sense.

- *Confidentiality* means that none of the nodes have been compromised. This definition is based on the assumption that all clients, servers, or other nodes under consideration contain or have access to critical, possibly the same, information. Therefore, if one node is compromised, that critical information is available to the attacker and so confidentiality of the overall system has been breached.

- *Integrity* means that a majority of the nodes (components) have not been compromised so that, if we request information from the components and compare results, at least one half of the results will match. Once an attacker has compromised more than one half of the components, we no longer have any confidence that the information being provided by a majority is correct. Byzantine failures (Lamport et al., 1982; tinyurl.com/klewe3x) can also be modelled in this framework whereby at least one third, a different but constant fraction, of the components need to be compromised for an integrity attack.

The time-to-compromise, t_i, of the ith node is a random variable distributed according to a probability density function, $f_i(t)$. The concept of time-to-compromise, discussed in more detail below, is based on the premise that any node is ultimately compromisable and the time when an attacker achieves the compromise is a random variable (which can include the attacker's skill level, resources, choice of attack strategies, and so on).

For example, the time to achieve success in a brute-force attack on a password would be distributed according to a uniform density between time 0 (when the attack begins) and time N/M where there are N possible passwords and M random passwords are tried per time unit. Techniques for estimating f_i and t_i for more complex computing systems have been developed and evaluated by the authors (Carin et al., 2008; tinyurl.com/mfyxu9r). Moreover, estimates of the time-to-compromise density allow us to estimate the cost-to-compromise of the ith component as well as the overall system or mission.

For simplicity, we assume that each density has the same form for each component and define α to be the lower bound on the support of f, β to be the upper bound on the support of f, μ to be the mean, and m to be the median. Moreover, we let n denote the degree of diversity (i.e., the number of distinct versions, for example, where clearly $n = 1$ represents a monoculture) as well as the number of parallel attackers such as would occur in a coordinated nation-state or organized crime attack.

Table 1 summarizes several analyses we have performed. The columns labelled Attackers and Diversity are as described; the entries in the columns for C, I, and A are the expected times to compromise confidentiality, integrity, and availability respectively. A graphical depiction of this analysis for the last line in the table is shown in Figure 1 to illustrate the wide variability on time-to-

Table 1. Expected times for an attacker to achieve one of the CIA goals

Attackers	Diversity	*C*	*I*	*A*
1	1	μ	μ	μ
1	n	μ	$n\mu/2$	$n\mu$
n	1	α	α	α
n	n	α	m	β

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

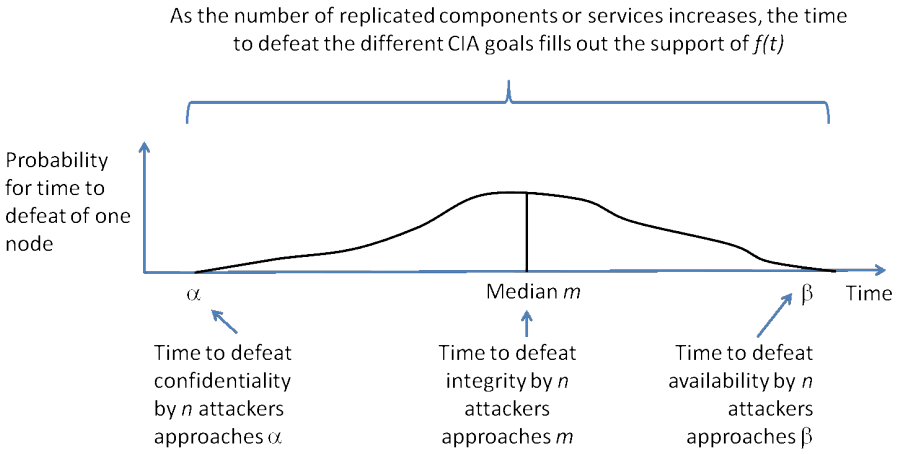

Figure 1. Expected times to achieve CIA security goals with n-fold diversity and n parallel attackers

compromise under these different scenarios. In this situation, where we have n-fold diversity and n parallel attackers, the expected times to achieve the various CIA security goals varies significantly. Decision makers must understand which security properties are most important to their organization's missions and invest accordingly.

This work quantitatively develops the trade space between confidentiality, integrity, and availability as a function of network diversity and time-to-compromise. In any such trade space, the IT manager must determine the "operating point of the design" or the balance between security properties and other important system properties such as "maintainability" and "mission utility".

It is informative to craft a simple opportunity-cost comparison based on this trade space. For instance, "cost-to-disrupt" is a cost to the adversary to compromise the enterprise that is directly estimated from the time-to-compromise scenarios provided above. This cost can be contrasted to the "cost-of-mission-disruption", which is a cost to the IT manager when considering the three types of security objectives (i.e., CIA) and the compromise of which disrupts the enterprise's mission (e.g., a disruption cost can be proportional to the number of users affected). Hence, analytically describing the trade space enables a richer strategic analysis regarding various IT enterprise objectives. This type of analysis is more explicitly described using the methodology in the next section.

Quantifying Cybersecurity Risk: The QuERIES Technique

The discussion above provides a framework and methodology for identifying various security goals and understanding the possible tradeoffs between them. Moving forward, if we are given a collection of identified security vulnerabilities impeding the achievement of the goals and possible defences or responses to those vulnerabilities, then we would next like to have some sense of how effective the proposed defences are in terms of performance metrics that go beyond a simple checklist.

In physical security, the time-to-compromise of a system is an accepted and measurable performance metric. Consider for example the case of the Overly Door Company (tinyurl.com/kdfdjm2), a supplier of US government General Services Administration approved Class 5 security vault doors suitable for storing national security information. These doors must provide protection against unauthorized entry for the following periods of time:

- 20 man-hours surreptitious entry
- 30 man-minutes covert entry
- 10 man-minutes forced entry
- 20 man-hours against manipulation of the lock
- 20 man-hours against radiological attack

Note that different times are specified for different types of attacks. Surreptitious entry means a method of entry that would not be detectable during normal use or during inspection by a qualified person. Covert entry

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

means a method of entry that would leave evidence, but would not be detectable by a user during normal use of the door and would only be detectable during inspection by a qualified person. Forced entry means a method of entry that would leave evidence of the attack and would be readily discernible in the normal use of the door; the attacker has no concern over leaving evidence that the vault door has been penetrated. Manipulation of the lock is defined as the opening of the combination lock without alteration of the physical structure or disarranging of parts. Ordinarily, manipulation would be accomplished by movement of the lock dial. Entry by radiological attack means the use of radioactive isotopes and other sources judged to be effective in determining the locks combination. Any entry made under these conditions within 20 man-hours shall be considered a failure of the vault door.

Physical security is relatively mature with much operational experience, so measures such as these have emerged as accepted standards within that community. Cyber-physical system security is still relatively new, so such performance measures are not yet standardized.

The authors have developed a technique, called QuERIES, for quantifying cybersecurity risk using ideas from a variety of disciplines and have demonstrated those techniques in software protection scenarios.

An example result of using the QuERIES methodology is depicted in Figure 2, which shows the distribution of times for completing a successful attack against a protected software system. The horizontal axis is time (equatable to cost in man-hours) and the vertical axis is the percentage of attempts requiring the corresponding time. The empirical probability density function was estimated using the QuERIES methodology and is depicted by the vertical bars. The specific compromise that was modeled in this example was an attack against a protected software system by an adversary whose goal is to extract specific parameter values from the protected code.

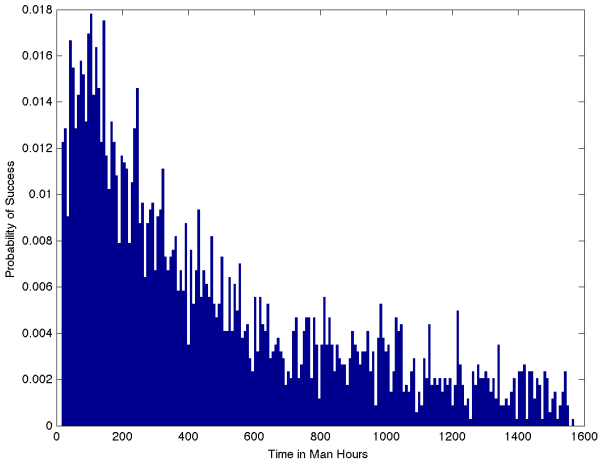

Figure 2. An example time-to-compromise density function for a software-protection defence developed by the authors in previous work on QuERIES

The plot in Figure 2 shows a probability density function for the time required by an attacker to compromise the protections. We model the time-to-compromise as a random variable in QuERIES because it depends on the skill level and approach an attacker takes. It might also depend on luck. Consider, for example, that we do in fact model brute-force password attacks in this way already – an attacker can be lucky and very quickly guess correctly but with very small probability. For a brute-force password attack, the corresponding plot would be a horizontal line at a very small probability going very far into the future.

The QuERIES methodology consists of a number of steps and has been successfully applied in a variety of cyber security situations. All seven steps in the QuERIES methodology are depicted on the right side of Figure 3; the four major ingredients of the methodology are shown on the left side of Figure 3 and are described below:

1. **Model the Problem**: Obtain objective quantities such as the economic value of the intellectual property (IP) (i.e., the protected software asset) to the IP owner; the cost of developing the IP by an adversary; the cost of obtaining the IP through other possible means; and a map of the specific protections applied to the IP asset.

2. **Model the Attacks:** Use the protection map and knowledge of reverse-engineering methodologies to build an attack graph represented as a partially observable Markov decision process (POMDP) (Russell and Norvig, 2002; tinyurl.com/lcpmldm).

3. **Quantify the Models:** Perform a controlled red-team attack against the protected IP and use another red- or black-hat team to conduct an information market for estimating the parameters of the POMDP.

4. **Use the Results:** The resulting estimates can be used to decide if the proposed protections are appropriate for the specific vulnerabilities in terms of various possible cost-benefit analyses.

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

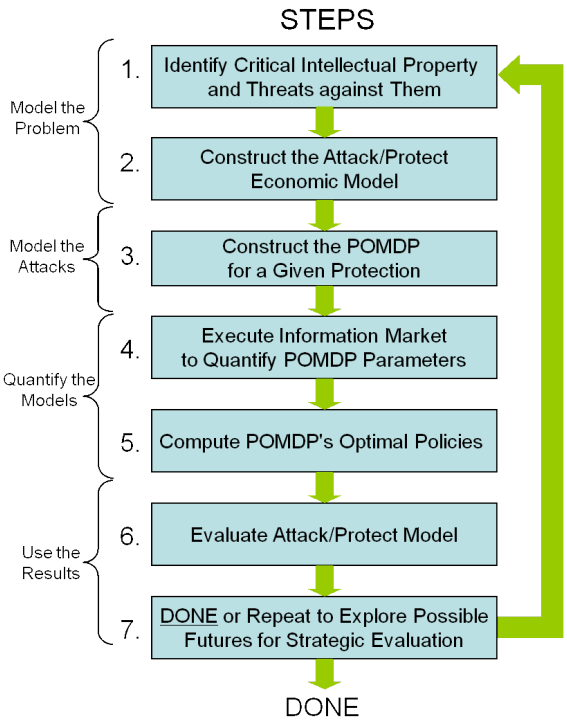

Figure 3. The seven steps of the QuERIES methodology

To illustrate such a cost-benefit analysis, consider the plot shown in Figure 4, which compares two approaches for attackers to decide when to stop an attack. The red line is the difference between the cost of the attack up to the corresponding time and the benefit of a successful attack; the blue line is the "cost-to-go" value of continuing the attack given that it has failed up to that time. The cost-to-go value is computed using dynamic programming based on the probabilities shown in Figure 2 and is similar to the techniques used for American Options pricing.

Figure 4 illustrates that a binary metric (i.e., true or false) is not suitable for determining whether or not a cyber-physical system can be compromised. Any system requiring a password can be compromised by a lucky guess and so would be considered insecure if that were the metric. Instead, we argue that the right kind of metric is, for example, the expected cost of a successful attack. If that cost is high enough, an attacker would not undertake the attack in the first place. This is the basis for all state-of-the-art encryption schemes, so our position on this is entirely consistent with existing practice and experience in that realm.

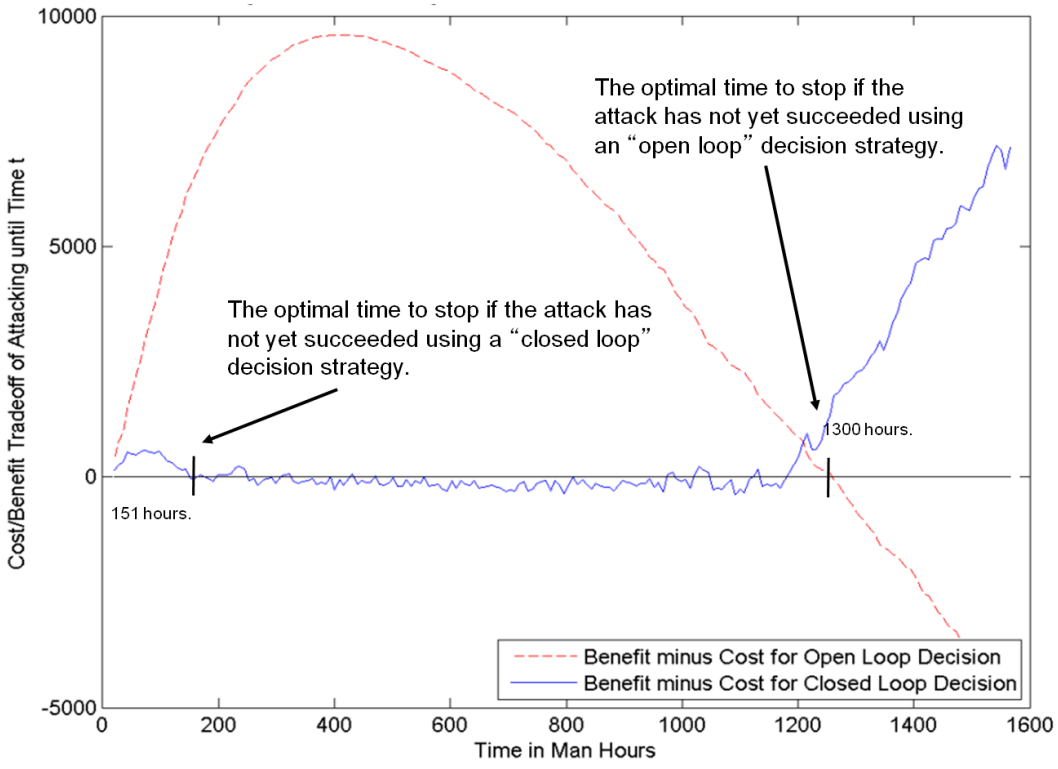

Figure 4. A comparison of two approaches to determine the optimal time for an attacker to stop an attack

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

The Three Tenets of Cybersecurity

Early in our work on threat and vulnerability analysis, we sought to identify simple – but still operationally meaningful, necessary, and sufficient – conditions for cyber-physical system vulnerabilities to exist. Once such conditions are identified, specific mitigations could be identified and evaluated.

This led us to identify three elements as being necessary and sufficient for successful attacks to occur:

1. The existence of inherent system susceptibilities
2. The threat's access to the susceptibility
3. The threat's capability to exploit the susceptibility

It is evident that, when these three elements are present in a system, an actual vulnerability exists.

Murphy's Law – "Anything that can go wrong, will go wrong." (Bell, 1989; tinyurl.com/llaps5q) – suggests that a system with vulnerabilities will be exploited given the appropriate operational environment. Moreover, a threat model that supports reasoning about whether an inherent system weakness rises to the level of a vulnerability is essential for cost-effective system-security engineering. This aspect is important because security for security's sake is neither affordable nor desirable, and so vulnerabilities must be quantified and only mitigated to the degree necessary to prosecute the enterprise's business processes or other missions.

We briefly describe these three elements below:

1. **System Susceptibility:** Absolute system confidentiality, integrity, and availability cannot be simultaneously achieved. Therefore, all systems will have design trade-offs resulting in inherent weaknesses. Such weaknesses will be manifest as software errors/bugs, protocol flaws, misconfigurations, or physical implementation constraints, and can be organized into the following eight categories of susceptibilities (National Vulnerabilities Database; nvd.nist.gov):

   a. *Input Validation Error (IVE):* includes failure to verify the incorrect input and read/write involving an invalid memory address. This category of susceptibility is also known as a boundary condition error (BCE) or buffer overflow (BOF).

   b. *Access Validation Error (AVE)*: causes failure in enforcing the correct privilege for a user.

   c. *Exceptional Condition Error Handling (ECEH):* arises due to failures in responding to unexpected data or conditions.

   d. *Environmental Error (EE):* triggered by specific conditions of the computational environment.

   e. *Configuration Error (CE):* results from improper system settings.

   f. *Race Condition Error (RC):* caused by the improper serialization of the sequences of processes.

   g. *Design Error (DE):* caused by improper design of the software structure.

   h. *Others:* includes susceptibilities that do not belong to the types listed above. This category of susceptibility is sometimes referred to as nonstandard.

2. **Threat Accessibility:** A threat will probe and analyze a system in order to discover which susceptibilities are accessible and how, with the goal of subsequent exploitation. Generally, the threat will use access points or services offered by a system to legitimate users as the original point of entry. Threat access is typically a superset of legitimate user access, because some access points may be undocumented or not of interest to legitimate users. Possible access points include wireless networks, legacy dialup lines, maintenance/service ports, automatic updates, and so on. Moreover, commercial and open source systems are accessible by the attacker for testing and exploit validation prior to launching a real attack. Any access offered an attacker provides a learning opportunity.

3. **Threat Capability:** After thorough surveillance (either via remote observations or in situ instrumentation) of the system design and its operation, an attacker will attempt to gain control, tamper with, or steal detailed system-design knowledge or other valuable data. Such attempts are often made using either a known or zero-day exploit determined after additional system reverse engineering. Skilled attackers typically employ a methodical approach to reverse engineering during which they expect to observe certain system behaviours. These system behaviours serve as exploitation guideposts and significantly aid

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

the attacker. The degree to which the attacker is successful will depend on their level of system knowledge, their ability and resources to develop and use specialized tools given the system's functionality and operating environment, and their overall level of reverse-engineering experience.

This form of threat model has deep roots in the Electronic Warfare (EW; tinyurl.com/kzrbp) test and evaluation community. That community shares a similar adversarial framework (i.e., measure-countermeasure) with cyber-physical system security. A version of this threat model was suggested for EW vulnerability analysis in 1978, and is called data link vulnerability analysis (DVAL) (Guzie, 2000; tinyurl.com/n4ge8w7). DVAL has four components in its vulnerability definition: susceptibility, interceptibility, accessibility, and feasibility. However, in contrast to DVAL, The Three Tenets threat model assumes that "feasibility" and "interceptibility" are effectively merged into what we call "capability." In today's complex cyber-physical systems based on commercial-off-the-shelf technologies, attackers can rehearse for almost any given operating environment and develop an exploitation capability. Such rehearsals are even possible with specialized computer-controlled systems, as demonstrated, for example, by Stuxnet (tinyurl.com/3vol5nk).

Thus, The Three Tenets threat model posits that three ingredients are necessary and sufficient for cyber-physical vulnerabilities to exist: i) a system susceptibility, ii) threat accessibility, and iii) threat capability. The three threat-model elements are illustrated in Figure 5. This figure depicts the co-occurrence of those ingredients as the space of vulnerabilities and therefore successful attacks.

Additional evidence that this vulnerability model is suitable comes from so-called routine activity theory (RAT) (Cohen and Felson, 1979; tinyurl.com/pml7vcq) that is used in criminology. This theory posits that crimes occur when three elements coincide: i) there is a motivated offender, ii) there is a lack of guardianship, and iii) there is a suitable target. The elements of RAT and the threat-model elements listed above have a clear correspondence: capability = motivated offender, accessibility = lack of guardians, susceptibility = suitable target. Our threat model is also related to "means, motive, and opportunity" arguments for convincing a jury of a suspect's guilt. The point is that previous work in criminology is relevant and consistent with our approach to cyberthreat modeling.

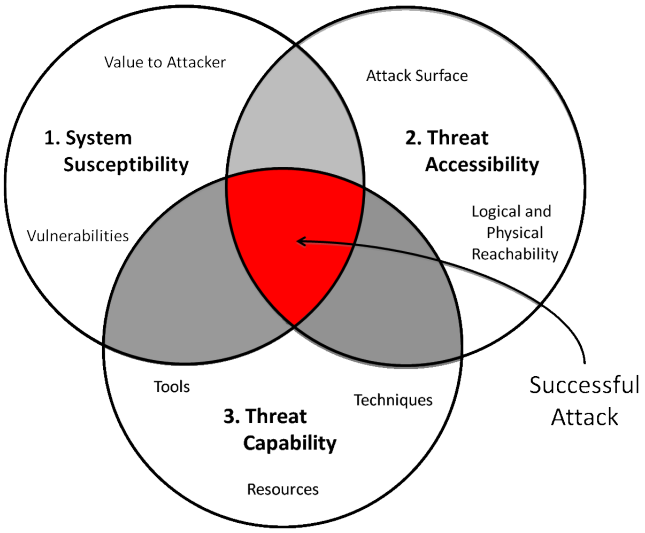

Figure 5. The three ingredients necessary and sufficient for cyber-physical vulnerabilities to exist

With these three necessary and sufficient conditions for cyber-physical vulnerabilities to exist, we can develop mitigation techniques and metrics for each condition. These mitigation techniques are called The Three Tenets and correspond to each condition outlined above. Collectively, The Three Tenets comprise a system security engineering approach consisting of both a secure design methodology and an assessment tool for security evaluation. The Three Tenets are introduced and described below:

1. **Focus on What is Critical:** The first Tenet instructs the designer to consciously and methodically focus on including only those system functions that are essential to the mission. This is an acknowledgement of Occam's razor (tinyurl.com/gxvu2) by the system designer. Adherence to this Tenet reduces the number of potential susceptibilities, and therefore, the paths between the attackers' starting state (i.e., the system access points) and goal states in which mission-essential functions, critical security controls, or critical data are compromised. This Tenet eliminates those access points and susceptibilities associated with unneeded functionality.

2. **Move Key Assets Out-of-Band:** The second Tenet instructs the designer to consciously differentiate between user access and attacker access for a given system's mission. This Tenet modifies system availability and is accomplished by moving the data/processes used by mission-essential functions, their

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

security controls, and associated access points out-of-band of the attacker either logically, physically, or both. By "out-of-band" we mean not accessible by the attacker through their preferred or available access methods. Adherence to this Tenet reduces threat access for a given mission (i.e., use case) and may enable unalterable observations of system state by a security control sensor. The extent and strength of access differentiation between the user and attacker is greatly influenced by the type of out-of-band mechanism employed and whether it is done in software or hardware.

3. **Detect, React, Adapt:** The third Tenet instructs the designer to employ dynamic sensing and response technologies (e.g., a security control sensor or reference monitor) that mitigate the threat's capabilities and exploitation attempts through automated (preferably autonomic) system behaviour. Adherence to this Tenet confounds the attacker's capabilities by making the system's defences unpredictable (i.e., nonstationary) and adaptive (i.e., with penalties) instead of merely being passive.

Just as each ingredient of the threat model has grounding in EW and classical criminology theory, each of The Three Tenets has been advocated and practiced in one form or another by computer security researchers and developers in the past. Further details and a more comprehensive treatment of The Three Tenets is available in a longer and more technical article (Hughes and Cybenko, 2013; tinyurl.com/l5wl5nt).

The Three Tenets provide a quantitative basis for the following security metrics, which are merely illustrative of more comprehensive and quantitative metrics that are possible:

1. **System Susceptibility Metric:** In its simplest instance, this system-construction metric instructs us to minimize the number of functionalities and services that act as access points to system-critical functions. This “reachability” metric is a direct consequence of the first Tenet: to identify, implement, and protect only what is mission critical.

2. **Access Point Metric:** Minimize the amount of input/output and system processes visible to an attacker. This metric is a direct consequence of the second Tenet: to move critical data “out-of-band.” Enumeration of “in-band” versus “out-of-band” access points is one way to measure application of the second Tenet.

3. **Threat Capability Metric:** Minimize useful insight into system operations in the sense that data observed at one time may or may not be similar or consistent with data observed at another time or on another system by the attacker. This “evidence variability” metric is a direct consequence of the third Tenet: to detect, react, and adapt. It is referred to by cybersecurity practitioners as a “moving target defence.”

These metrics can be readily measured by an enumeration of access points and data input/output or process observations together with determination of system functional behaviour.

Moreover, there are clear economic and effectiveness tradeoffs between, for example, implementing Tenet 3 (detect, react and adapt) and Tenet 1 (implementing only mission-critical functionality). These tradeoffs can be addressed through a QuERIES-type methodology and are the subject of ongoing work.

Conclusion

In this article, we have presented threat-driven, descriptive security methodologies that enable reasoning about cyber-physical system design in a strategic fashion. We feel that this approach is a clear alternative to traditional prescriptive approaches to cybersecurity that provide little insight into the comparative value of security solutions given the entirety of the system security trade-space. Underpinning our methodologies is the concept of “time-to-compromise.” We suggest that this is a fundamental metric associated with any adversarial environment and that cyber-physical system security is no different than physical security in this respect. Concrete metrics are described that are functionally related to and expand upon time-to-compromise. These metrics serve as informative and quantitative guides in secure system design. Future work will describe the mathematical underpinnings of The Three Tenets and provide a more complete derivation of the resultant quantitative security metrics. Additionally, the benefits of analyzing complex system security by employing probabilistic formulations such as QuERIES and the CIA analysis will be illustrated via reduction to practice for varying use cases. Finally, we intend to develop a more explicit coupling of these methodologies to a lifecycle security analysis for cyber-physical systems.

Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity

Jeff Hughes and George Cybenko

About the Authors

Jeff A. Hughes is President of Tenet 3, LLC. Tenet 3 is a cybertechnology company with a focus on autonomous cyber-physical systems, analyzing their trustworthiness, and evaluating economical ways to demonstrably mitigate security risks. Previously, Jeff held various positions in the US Air Force Research Laboratory (AFRL), where he led research into advanced techniques for developing and screening trustworthy microelectronic components and performing complex system vulnerability and risk analysis for cyber-physical systems. Jeff has an MS in Electrical Engineering from the Ohio State University and has completed graduate work towards a PhD at the Air Force Institute of Technology in Ohio, United States.

George Cybenko is the Dorothy and Walter Gramm Professor of Engineering at Dartmouth College in New Hampshire, United States. Professor Cybenko has made multiple research contributions in signal processing, neural computing, information security, and computational behavioural analysis. He was the Founding Editor-in-Chief of both *IEEE/AIP Computing in Science and Engineering* and *IEEE Security & Privacy*. He has served on the Defense Science Board (2008-2009), on the US Air Force Scientific Advisory Board (2012-2015), and on review and advisory panels for DARPA, IDA, and Lawrence Livermore National Laboratory. Professor Cybenko is a Fellow of the IEEE and received his BS (Toronto) and PhD (Princeton) degrees in Mathematics.

Citation: Hughes, J. and G. Cybenko. 2013. Quantitative Metrics and Risk Assessment: The Three Tenets Model of Cybersecurity. *Technology Innovation Management Review*. August 2013: 15–24.

Keywords: quantitative cybersecurity, risk assessment, vulnerabilities, confidentiality, integrity, availability