Evaluating Security Controls Based on Key Performance Indicators and Stakeholder Mission

F. T. Sheldon, R. K. Abercrombie
Oak Ridge National Laboratory
1 Bethel Valley Road, PO Box 2008
Oak Ridge, TN 37831-6418, USA
+1 865-576-1339, +1 865-241-6537
sheldtonft@ornl.gov, abercrombier@ornl.gov

A. Mili
Department of Computer Science
New Jersey Institute of Technology
Newark, NJ 07102-1982, USA
+1 973-596-5215
ali.mili@njit.edu

Abstract
Good security metrics are required to make good decisions about how to design security counter-measures, to choose between alternative security architectures, and to improve security during operations. Therefore, in essence, measurement can be viewed as a decision aid. The lack of sound practical security metrics is severely hampering progress in the development of secure systems.

The Cyberspace Security Econometrics System (CSES) offers the following advantages over traditional measurement systems: (1) CSES reflects the variances that exist amongst different stakeholders of the same system. Different stakeholders will typically attach different stakes to the same requirement or service (e.g., a service may be provided by an information technology system or process control system, etc.). (2) For a given stakeholder, CSES reflects the variance that may exist among the stakes she/he attaches to meeting each requirement. The same stakeholder may attach different stakes to satisfying different requirements within the overall system specification. (3) For a given compound specification (e.g., combination(s) of commercial off the shelf software and/or hardware), CSES reflects the variance that may exist amongst the levels of verification and validation (i.e., certification) performed on components of the specification. The certification activity may produce higher levels of assurance across different components of the specification than others. Consequently, this paper introduces the basis, objectives and capabilities for the CSES including inputs/outputs and the basic structural and mathematical underpinnings.

This manuscript has been authored by a contractor of the U.S. Government under contract DE-AC05-00OR22725. Accordingly, the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
CSIIRW'08, May 12–14, 2008, Oak Ridge, Tennessee, USA
Copyright 2008 ACM 1-60558-098-2…$5.00.

Measurement Criteria
Some qualities of a good metric include: (1) Ability to measure the right thing (e.g., supports the decisions that need to be made), (2) Quantitatively measurable (e.g., damages in dollars of profit loss), (3) Capability to be measured accurately, (4) Ability to be validated against ground truth, and (5) Confidence level one has in the assertions made within the framework of the metric. To these criteria, one should add the following desirable properties: (1) Inexpensive in time and cost to execute, (2) ability to be refereed independently, (3) repeatable so the outputs are independent of the analyst performing the measurement, and (4) scalable from small, single computer systems to large, nation-scale enterprise networks.

MFC Characterization
System security should be characterized, not by some abstract discrete scale, but rather by the very concrete determinant, mean failure cost (MFC) [1]. MFC reveals how much each stakeholder stands to lose from mission value due to the lack of security. Subsequent use of this quantification enables us to derive an economic model that captures the tradeoffs involved in deploying security counter measures. The CSES measurement process proceeds in three steps (Generation of Stake Matrix, Dependency Matrix, and Impact Matrix), discussed in the subsequent subsections [2]. Let us consider the fundamental elements needed for evaluating security controls.

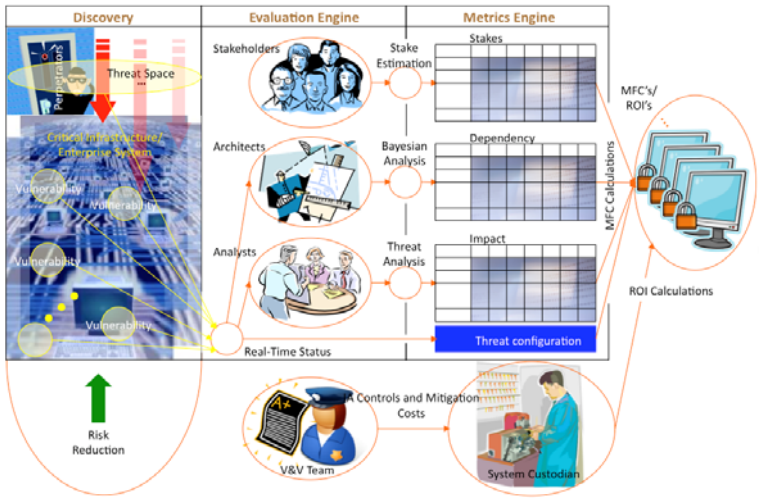

Figure 1. Cyber Security Econometrics (CSE)

Fundamentals

Figure 1 shows essential I/O components and phases (i.e., discovery, evaluation and metrics) including data collection/analysis and consists of the following parts: *System Stakeholders* refers to any person or organization that has a stake in the operation of the system (i.e., users, operators of the system, hosts of the systems, etc.); *Security Specification* is used in the same way that correctness is a relative attribute (i.e., a system is correct with respect to it's functional specification) and refers to a representation of the security attributes that a system must satisfy to be deemed secure; *Security Requirement* is used in the same way that a complex functional specification is typically composed of simpler components (i.e., representing elementary functional properties), and is also more generally composed of simpler security requirements; *Mean Failure Cost* is used in an operational sense because the lack of security within the system may cause damage, in terms of lost productivity, lost business, lost data, resulting in security violations. We represent this loss by a random variable, and define MFC as the mean of this random variable [1]. As discussed further, this quantity is not intrinsic to the system, but varies by stakeholder.

Generation of Stakes Matrix

The MFC estimation depends on the following principles: (1) a stakeholder may have different stakes in different security requirements, and (2) a security requirement may carry different stakes for different stakeholders. The best way to represent these dependencies is through a 2D matrix, where the rows represent stakeholders S_i, the columns represent security requirements R_j and the entries represent stakes (see Table 1). The FC entry at row i, column j, represents the cost that stakeholder S_i would lose if the system failed to meet the security requirement R_j (i.e., also represented as $FC(S_i,R_j)$). Each row is filled by the corresponding stakeholder, possibly in their own (possibly distinct) financial / economic terms (Dollars, Person Months, Euros, etc).

The estimation of MFC presumes that 1) the same stakeholder may have different stakes in different security requirements, and 2) the same security requirement may carry different stakes for different stakeholders.

Generation of Dependency Matrix

The *Dependency Matrix* is used to estimate the probability that a particular security requirement is violated in the course of operating the system for some period of time. In this way, we link the probability of failing a particular requirement with the probability of failure of a component in the system. The elucidation of this probabilistic link involves an analysis of the system's architecture, to determine *which component contributes to meeting which requirement.*

The analysis of the system architecture, by architecture subject matter experts, enables the derivation of conditional probabilities that link the probability of component failures with the probabilities of failing to meet specific requirements. The term $\pi(E_j)$ represents the probability of event E_j and the term $\pi(R|E_j)$ represents the probability of

Table 1. Stakes (ST) Matrix: Cost of failing a security requirement.

		Security Requirements				
		R1	**R2**	**R3**	**…**	**Rn**
Stakeholders	**S1**					
	S2					
	S3					
	…				FC_i^j	
	Sm					

failing to satisfy requirement R_i, given the hypothesis E_j (i.e., that event j has occurred). Table 2 assumes there exists a component event E_j for a requirement R_i where the probability of failure to satisfy requirement R exists ($\pi(R_i \mid E_j)$).

The probability of failing a requirement is obtained by the sum, for all components, of the conditional probabilities of failing that requirement, conditional on failure of the component, weighted by the probability of failure of the component. The conditional probabilities are given by the dependency matrix. Component failure probabilities are the subject of the next section.

Table 2. Dependency (DP) Matrix: Links requirements and components.

		Components				
		C1	**C2**	**C3**	**…**	**Ck**
Requirements	**R1**					
	R2					
	R3					
	…				$\pi(R_i \mid E_j)$	
	Rn					

Generation of Impact Matrix

The Impact Matrix shows the Component Failure versus Threats Relationship Grouping (Table 3). The probability of component failure depends on three factors: (1) Technical Controls that the component is provided with to protect against threats and to mitigate damage in the case of successful attacks, (2) Pattern of threats that the component is subjected to, and (3) the degree of verification and validation that the component has undergone.

To assess the likelihood that a particular threat leads to the failure of a component, we consider a set of cataloged threats (or families of threats with common attributes), say *T1, T2, T3, … Th*, and we consider the events *V1, V2, V3, … Vh, Vh+1*, where *Vi*, for $1 \leqslant i \leqslant h$, stands for: *Threat i has materialized*, and *Vh+1* stands for: *No threat i has materialized*.

This equation links the probability of threat *Tj* (which is $\pi(Vj)$) to the probability of component failure for component *Ci* (*which is* $\pi(Ei)$). To apply this formula, we need to derive the conditional probabilities, which we propose to represent in a 2 dimensional matrix, that we call the ***Impact*** matrix.

Conclusions

We summarize our findings:

• The vector of mean failure costs (*MFC*, one entry per stakeholder) is given by the following equation:

$$MFC = ST \circ PR,$$

where *ST* is the ***Stakes*** matrix and *PE* is the vector of requirement failure probabilities (simply one entry per requirement).

• The vector of requirement failure probabilities is given by the following equation:

$$PR = DP \circ PE,$$

where *DP* is the ***Dependency*** matrix and *PE* is the vector of component failure probabilities (simply one entry per component).

• The vector of components failure probabilities is given by the following equation:

$$PE = IM \circ PV,$$

where *IM* is the ***Impact*** matrix and *PV* is the vector of threat emergence probabilities (simply one entry by type of threat).

By substitution, we find the equation that gives us the vector of mean failure costs of all stakeholders as:

Table 3. Impact (IM) Matrix: Links threat relationship groupings.

		Threats				
		T1	**T2**	**T3**	**…**	**Th**
Components	**C1**					
	C2					
	C3					
	…				$\pi(E_i \mid V_j)$	
	Ck					

$$MFC = ST \circ DP \circ IM \circ PV.$$

Utilizing a user interface, the stakes matrix (*ST*) is filled by stakeholders according to the stake they have that individual requirements are satisfied (i.e., assurance); the dependency matrix (*DP*) is filled in by the system architect according to how each component is affected by each threat.

The remaining question is: how to fill the vector PV, that represents the probability of emergence of the various threats that are under consideration? This is done empirically, by simulating and/or operating the system for some length of time and estimating the number of threats that have emerged during that time and continue to be refined as the system evolves. From these numbers, we infer the probability of emergence of all the threats during one hour of operation.

REFERENCES

[1] A. Mili and F. T. Sheldon, "Measuring Reliability as a Mean Failure Cost," in *Proceedings of the 10th IEEE High Assurance Systems Engineering Symposium* Dallas, TX: IEEE, 2007, pp. 403-404.

[2] R. K. Abercrombie, F. T. Sheldon, and A. Mili, "System and Method for Implementing and Monitoring a Cyberspace Security Econometrics System and Other Complex Systems," in *US Patent Office,* US Patent Pending, May 12, 2008, pp. 1-95.