

12th ICCRTS
“Adapting C2 to the 21st Century”

Security Metrics for Communication Systems

Communications System Security
Security Metrics
Metrics and Assessments

Mark D. Torgerson
Cryptography and Information Systems Surety Department

Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-0785
505-284-5677 or 435-843-7283
mdtorge@sandia.gov

Abstract

This report discusses the possibility of creating meaningful security metrics for communication systems. In particular, we examine security metrics from an axiomatic standpoint and prove that it is not possible to measure trust in an absolute sense. We do not conclude that it is impossible to create a secure communication system; rather we argue that it is impossible to detect the occurrence. We also explore directions where further research is possible.

DRAFT

This Page Intentionally Left Blank

1. Introduction

There is an ongoing evolution in the communications industry. Experts are becoming more keenly aware that communications systems are valuable commodities that are increasingly under attack. In an effort to reduce the number of successful attacks and thus stem the tide of loss associated with a compromised communication system, security experts are employed to build defenses around or within the communication system to prevent adversarial manipulations of that system. The security community has matured to the point that ad-hoc methods of security and evaluation are deemed insufficient. Thus the community seeks some methodology that would allow systems to be evaluated against and given some score or metric as to the level of security of the system.

In November, 2005 the Infosec Research Council (IRC) published a document entitled *Hard Problem List*[1]. That document delineates a number of outstanding hard problems that the security community needs to solve. One of the hard problems presented by the IRC is that of coming up with enterprise level security metrics. They also identify several areas or subcategories that range from definitions to composability.

The intuitive notion of a security metric is that of a function or process where one would input a communication system and out would come a number, or set of numbers, indicating the level of security of the system. The intent of this paper is to discuss the possibility of creating a security metric that meets that intuitive notion yet has some rigor associated with it. We will argue using axiomatic reasoning that it is not possible to define security metrics that fit with the intuitive notion of such metrics. Further we argue that the metrics one is, potentially, able to define are only of limited value in measuring security.

2. Initial Details

We first give some notional descriptions of terms used in this paper.

Communications System (System): A real collection of hardware, software, and human components brought together to facilitate communications of some kind.

Adversary: An entity that desires to gain some nefarious goal against the system.

Security subsystem: The system components used, either directly or indirectly, to prevent an adversary from achieving his goals.

Weakness: An attribute of a system that an adversary may use while attempting to achieve his nefarious goals.

Trust: Confidence that one may have in their system in preventing an adversary from achieving his nefarious goals.

One of the difficulties here is rigorously defining the terms above. It is clear that the descriptions above are not sufficient in many circumstances. However, further refinements always lead to lengthy descriptions that are no better defined, incomplete, or

contradictory. In fact, attempting to be absolutely precise about these definitions will always lead to logical quagmires of one kind or another.

In any logical system, the system must be based on a collection of undefined terms, terms that everyone knows or assumes the meaning of, have nice descriptions, but do not have precise definitions. Every student that has taken a class in mathematical logic has gone through an exercise of attempting to get to the root of some definition sequence. Assuming that there are no undefined terms, term A is defined by terms B_1 through B_k . Each of these terms is defined by other terms and so on. Eventually one will either begin to repeat terms or use term A to define some subsequent term. Neither type of circular definition is allowed, and thus the argument that every logical system requires the use of undefined terms.

In creating a logical system, one has some freedom in choosing those undefined terms. For instance, in Geometry a 'point' is a typical starting place. (Try defining a point.) One can describe a point, and one can spend hours trying to acquire a visual image of the concept. One can even acquire a very satisfying notional feeling for what a point is. However, any attempt to rigorously define a point always leads to very messy descriptions of other Geometric objects that can be constructed from points or which have no intuitive quality to them. One finds it difficult to attain the same level of satisfaction with the quality of those objects. Points are wonderfully described, but horribly defined and thus often left as an undefined term.

Once chosen, those undefined terms are then fixed in the system and should have certain properties. For instance, one undefined term should not be expressible by the others; in that case, it is not really an undefined term. Axioms are an extension of the undefined terms. They are generally understood to be correct notions or even defining properties of the system that cannot be proven or implied by previously described terms or axioms. In one of the great achievements in logical systems, Gödel was able to show that no system of axioms is at the same time complete and consistent.

The goal of this work is to look at the notion of trust metrics from an axiomatic viewpoint. Our slant is to take the terms given above and treat those as undefined terms in our logical system. (We all know or have an intuitive feel as to what an adversary is, but just try to define one rigorously, comprehensively, consistently, and in a way that does not brook argument from someone else.)

3. Adversary

In general there are two important dimensions to consider when discussing an adversary. The first is his knowledge and second his physical resources. The resources include computational ability, as well as other things such as money, having an ability to conduct side channel attacks, having the tools to pick locks, having corrupted an insider, time-travel, etc. A designer's understanding about an adversary's physical resources plays a key role in the designer's decision making process during the design of a security system. For instance, key length for symmetric key encryption algorithms is based on the

difficulty for an adversary to exhaustively search through the set of keys. At this point 128 bits is common, where most believe that that number of bits will be sufficient for quite some time in the future. However, a sufficiently strong quantum computer may make today's 128-bit keys obsolete. System designers know this and rely on the fact that we are many years from making a viable quantum computer. On another, similar, but far-out vein there may be a time when someone will invent a device that allows some sort of phase shift or teleportation. In that case, many of the physical protections placed around a communication system could be circumvented simply by walking through walls or just appearing in sensitive areas. If one were to let their imagination run away, it is clear that every communication system in existence today is completely trivial to some future adversary with the right physical resources.

Similarly, if a company whose total assets value in the few millions has an adversary who is willing to spend billions to retrieve certain company proprietary data, the adversary may as well just buy the company and own the data.

The discussion of metrics has to be scoped in a way that makes sense given adversaries (current and future) with significant resources in mind.

When a group of security experts get together and talk about secure communication systems, invariably someone brings up the point that no system is 100% secure. Sometimes that incites discussion about the meaning of security and so on, but few argue the intuitive meaning behind the statement. The notion is certainly true in the case of a completely resource unbounded adversary. On the other end of the spectrum, protecting against adversaries with extremely limited resources may be possible. As a boundary condition, it may be possible to create a security system capable of protecting against an adversary with zero resources. Even an unprotected system is safe from a person who is in a coma.

One may go quite far with very limited physical resources given the proper knowledge and opportunity. An adversary may walk past a napping guard, enter a door propped open by the guard for convenience in letting people in and out of an area deemed sensitive, and then read communications printed out earlier for someone coming to pick them up later. Or a child with limited programming skills may download a very potent root kit created by someone with large amounts of know-how and resources.

A real communication system of any value will have real adversaries, with non-zero but realistic resources. Unfortunately, it may not even be possible to identify one's adversaries. Even if those adversaries have been identified, it is nearly impossible to measure their current resources. Nor can one expect that a particular adversary will have the same resources tomorrow.

The ability of an adversary to gain greater knowledge about a system and its weaknesses is a real and immediate concern. Problems and bugs with communication products are found all the time. Each newly discovered problem will amount to, possibly, another

avenue for attack by an adversary. Sometimes those discovered weaknesses are only exploitable after some, possibly, significant change in resources.

In an attempt to scope the notion of an adversary to allow one to make sense of metrics, we will say that a particular adversary A is represented by (K,R) where K is the adversary's knowledge and R is his physical resources. Because, if we let both parameters be unbounded, every system is trivial and talking of metrics makes no sense whatsoever, we will assume throughout that each adversary discussed will have a fixed and bounded set of physical resources. That is, each adversary is able to learn new things about the system but is not able to raise their physical resources or capabilities past a certain point without becoming a different adversary. An adversary with resource bound B may be written $A=(K,R \leq B)$.

4. Weaknesses

It does not make sense to talk of system weaknesses without the context of the physical resources available to the adversary. As discussed above, every real communication system in use today is weak against the right adversary with the right physical resources. Without placing some bounds on the adversary, talking about a set of system weaknesses is meaningless. So, to provide a meaningful context and make the notation more palatable we assume that all adversaries have a fixed, nonzero, resource bound B . We tacitly assume that any discussion implicitly assumes and incorporates that resource bound.

Going back to the statement "No system is 100% secure," we translate into our first axiom.

Weakness Axiom 1: Every real communication system has a non empty set of weaknesses.

We assert that the set of weaknesses of a system may be very large, if not infinite. That set of weaknesses for a given system is fixed and is independent of the owner's or any particular adversary's knowledge of those weaknesses. It may be a debatable point as to the existence of a weakness before it is discovered. Encryption algorithms were not designed to resist linear cryptanalysis before linear cryptanalysis was discovered. After that discovery, certain existing encryption algorithms were shown to have a weakness toward that attack. The encryption algorithms did not change, rather our knowledge and understanding of their weaknesses changed. Further, even today, not all analysts and designers understand and can apply linear cryptanalysis with the same degree of facility.

Here, we take the stance that just because an analyst does or does not understand the power of a particular method of attack, does not affect the strength or weakness of a system against that method of attack. Even further, we feel that, in an existential sense, a system may indeed have undiscovered weaknesses to undiscovered attack methodologies.

We want to emphasize again that we are discussing a real communication system. There are specific algorithms that have no theoretical weaknesses. For instance a one-time pad

where key is generated from a true random source has the property of perfect confidentiality, even in a theoretical sense. However, the “on paper” algorithm is a far cry from a real communication system, where the implementation, the supporting processes, and human factors all come to bear on the security of the system implementing a one-time pad. Further the one-time pad does not give any sort of authentication of source or data integrity.

A given system S will have certain security protections placed on it. This security subsystem may be simple or extremely complicated and elaborate. The protections provided by the security subsystem are designed to make it difficult for an adversary to accomplish his goals. Certain weaknesses are mitigated by the protections.

In most cases, the security subsystem is an integral part of the overall system S and cannot be decoupled from it. Any changes to the security subsystem will most certainly change the system itself. Those changes will remove some weaknesses and will likely introduce new weaknesses. One may view the changed system as a completely new system with an entirely new set of weaknesses, we do not. For argument’s sake, we assume that the security subsystem can be identified and decoupled from the communication system and that changes to the security subsystem mitigate only existing weaknesses and do not introduce or create new weaknesses.

Let S be a given communication system. Let W be the set of weaknesses of S and let P be the protections placed on S . Let $MW(P)$ be the set of weaknesses mitigated by P . Similarly, let $UMW(P)$ be the weaknesses that are left unmitigated by P . When the context is clear, we suppress P in the notation. Note that MW and UMW are disjoint sets whose union is W . Note also that for a given set of protections (and adversarial resource bound), they are constants of the system and, like W , are not affected by who is or who is not inspecting the system, their knowledge or lack thereof.

For each viewer V of the system, let $WK(V)$ and $WUK(V)$ be the set of weaknesses that are, respectively, known and unknown to V . We have then that WK and WUK are disjoint sets whose union is W . The weaknesses exploitable by V are then

$$E(P,V) = UMW(P) \cap WK(V),$$

the unprotected weaknesses that are known to V . When P is understood we will write E_V to denote the weaknesses exploitable by V . One definition of a secure system would then be, “A system is *secure* against adversary A if E_A is the empty set.”

Also assume that in time, V will gain a better understanding of the system, its protections and its weaknesses. It is certainly the case that V may forget about previously known weaknesses and protections, however, for simplicity, we will assume no viewer forgets. Thus, in time, WK will grow in a set inclusive way.

The owner of the system has a special view of the system and is able to make certain changes to the system when weaknesses are discovered. There will always be a lag between when the owner discovers a weakness in the system and the time that he is able

to beef up the security subsystem to mitigate that weakness. One would hope that that lag is short, but it may be quite lengthy if the owner does not have the resources to adjust his fielded system. If O is the owner of the system, then when E_O is non-empty, the owner of the system knows of specific system weaknesses for an adversary with resource bound B that are not mitigated by the security system.

Weakness Axiom 2: For all viewers, V , of the system we have that $WK(V)$ is a strict subset of W .

It is certainly reasonable to assume that in the near (and distant) future, mankind will not know all there is to know about security and the building of secure communication systems. Someone may indeed discover a new type of attack against a particular system. This notion of discovery turns out to be one of the real culprits in the attempt to define reasonable security metrics. This axiom is particularly important when the viewer is the owner of the system. When the viewer is the owner, Weakness Axiom 2 says that the owner will never be able to identify all system weaknesses. This is a little different than Weakness Axiom 1 which just says that the weaknesses exist. Weakness Axiom 2 can be stated another equivalent way: For all viewers, V , of the system we have that $WUK(V)$ is a nonempty subset of W .

Weakness Axiom 3: The system owner cannot know $WK(A)$ for all adversarial viewers of the system.

An owner of the system likely does not know exactly who his adversaries are. Even with the adversaries that are known, the system owner may only conjecture as to what a particular adversary does or does not know about the system and the system's weaknesses. Knowledge beyond conjecture about a particular adversary may be gained through an active pursuit of that adversary.

Together these three axioms paint a bleak picture for the security world. Together they say that a system will always have weaknesses, and no matter what you do, you can never discover all the weaknesses or patch all the holes in your system, and you will never be certain of what your adversaries know about those holes or if the adversary knows something you do not.

5. Security Metrics

When people talk of security metrics, they often envision a function or device wherein a system can be input and out pops a goodness rating for the system. The goodness rating is to be independent of adversaries, real or imagined, their knowledge or their capabilities, present or future. That security metric is also wished to be meaningful in time and across comparison of other systems.

A communication system *metric* is a computable function from the set of communication systems into the real numbers. A *security metric* is a metric that gives some indication as

to how secure the system is. Here, computable means that a viewer of a given system can input the system (or aspects of the system) into the function and, in a real, deterministic way, evaluate the metric.

Desirable attributes of a security metric are that they should not be trivial and should be meaningful. After all, the purpose of the metric is to give the owner a certain level of confidence in the security of the system. There are several metrics that fail these desires. For instance, one could give a fixed value to every system. Even though attaching a rating of 42 to every system is a real-valued function of the system, the metric is trivial and useless. Similarly, converting the name of the system into a real number is also not likely useful or meaningful from a security standpoint.

The notion that a real metric be computable leads to a very practical requirement, which, when given, allows a rule of thumb that unknown sets cannot be deterministically processed or evaluated. This rule of thumb leads to another axiom.

Metric Axiom 1: The only metrics that evaluate unknown sets are trivial functions of the unknown sets.

From a functional standpoint one would like a security metric to be robust enough to process different systems in a way that final results would be comparable. However, it is not clear that it makes sense to attempt to compare the security of two completely different communication systems that have completely different functional requirements and operational goals. So, we restrict our discussion to the examination of the smaller problem of measuring the security of a family of communication systems that are based on the same underlying system but differ in some way.

5.1. Weakness Based Metrics

There are conceivably many different types of security metrics, however it is hard to conceive of a meaningful security metric that does not take into account system weaknesses. In an effort to get a handle on the more general case we will restrict our view in this section to metrics based solely on system weaknesses.

Here, the scenario that we would like to mimic is where the owner, O , of the system, S , would like to place protections on the system to mitigate the weaknesses in the system. If a security metric could be created, the system owner might then evaluate different possible mitigation strategies and determine the best path to satisfy his needs.

Let $WK = WK(O)$ be the fixed set of weakness known by the owner of S . Given O , S , W , and WK as described, let \mathbf{W} be the power set of W . For each $w \in \mathbf{W}$ let S_w be the communication system based on S with w being the set of mitigated weaknesses. There is a one to one correspondence between the communication systems defined by the mitigated weaknesses and the sets of mitigated weaknesses. Even though we are speaking of the systems, the bijection described will simplify the notation and discussion.

A *Weakness Based Metric* is a security metric wherein the owner, or designee, inspects the system, identifies the weaknesses that have been mitigated and assigns a real number to that system. That is, it is a computable function mapping \mathbf{W} into the real numbers \mathbb{R} , where we have $M:\mathbf{W}\rightarrow\mathbb{R}$.

Theorem 1: There are no weakness based metrics that include $WUK(O)$ in a non-trivial way.

Metric Axiom 1 and Theorem 1 are obvious, if not absurd. Yet the implications are extensive. A security metric that is independent of adversaries, real or imagined, present or future must measure the totality of the system weaknesses. The weaknesses unknown to the one doing the measuring are not measurable. Thus, a metric that tells one how secure their system is, in an absolute sense, not achievable.

A weakness based metric is *consistent* if for all $w \in \mathbf{W}$ we have $M(w)=M(w \cap WK)$.

A consistent weakness based metric is the best weakness metric that one can hope for. The system owner can include only the weakness that he knows about in his metric evaluations. The reason for the definition of a consistent metric is subtle and mainly for technical precision. It may be that one viewer of the system has a fairly comprehensive understanding of the system weaknesses. That viewer of the system may define a metric based on their view. That metric when restricted to the owner's view would not necessarily be consistent. However, from the owner's point of view, the sets $w \in \mathbf{W}$ may not actually make sense; his universe is always restricted to sets of the form $w \cap WK$.

Now suppose M is a valid consistent weakness based security metric. Suppose also that w is the set of migrated weaknesses and that $M(w)=42$. What could this number possibly mean? Can this number give any indication of how secure S_w is? The answer is a clear no. As defined above, the V exploitable weaknesses

$$E_V = \cup_{A \in \mathbf{A}} M(A) \cap WK(V)$$

are the only weaknesses that a system owner needs to worry about. Let $E = \cup_{A \in \mathbf{A}} E_A$ where the union is taken over all adversaries A . The set E is the collection of all exploitable weaknesses collected from all adversaries and may be considered to be the sum total of the weaknesses that are not protected against, that one has to worry about.

Theorem 2: The set E is not measurable by the owner of the system, and thus no metric exists which includes that quantity in a non-trivial way.

There are three sources of immeasurability. First, the set of unmitigated weaknesses is not measurable by anyone, let alone the owner of the system. Second, $WK(V)$ is not known by the owner of the system and thus is not measurable by the owner of the system. Thirdly, no system owner can know the totality of his adversaries.

With enough legwork, a system owner may identify a few adversaries and get a vague idea of their physical capabilities, but it is a difficult thing to get a true measure of the things that the adversary knows about, above and beyond what the system is able to

protect against. Even if for a particular adversary E_A were measurable, such a measure would not be a measure of the system, rather, it would be a measure of that particular adversary. One cannot look solely at the system and compute such a measure. In fact, one would likely not even input the system to get the measure, rather one would set the system aside and expend resources to better understand what the adversary knows or does not know about the particular system.

We argue that the best security metrics that one could hope for are those that measure the totality of the unmitigated weaknesses. The second best measure the collection of exploitable weaknesses. We have shown that no such metrics exist.

By ruling out the unknown quantities, we have argued that a metric on the system may, at best, be a measure of known quantities, such as MW and/or E_O , the mitigated weaknesses and the exploitable weaknesses that the owner knows about. In all cases, the metric can measure only the weaknesses that the system owner knows about. Unfortunately, measures based on those sets do not give a true picture as to how secure the system is, how resistant the system is to being exploited by an adversary. A naïve owner may give a large rating to a system that he is confident in, but yet the system is completely vulnerable to certain attacks that he is unaware of.

5.2. Protection Based Metrics

Even though the most robust measures of security are not available, one must still attempt to quantify the quality and security of one's system. We now focus on the system protections themselves.

Basing one's metrics on protection systems is something that may be done. For instance, the cost to implement is one measure. Within reason, this does give a vague notion of how secure a system is, in a relative sense. Another possibility is that of using certification levels defined by an outside party. All of these are processes that one may run the protections through and get some relative sense of how well protected a system is.

For a given system, S , let \mathbf{P} be the set of possible protection systems that may be placed on S . *Protection based metrics* are those security metrics that are a function of the space of possible protections of a system. Note that for each $P \in \mathbf{P}$ there is an associated set of mitigated weaknesses. Thus there is a mapping $Z: \mathbf{P} \rightarrow \mathbf{W}$.

The mapping Z is potentially many to one. Different protection systems may mitigate the same set of weaknesses. For instance, competing vendors may design similar protection systems that mitigate the same weaknesses; however, they may differ in efficiency and maintainability or other less security related aspects such as price, ease of use, etc.

The fact that two different sets of protections may lead to the same set of mitigated weaknesses means that there may be metrics on the protection space that do not correspond directly to a metric on the set of weaknesses. On the other hand, one may

compose a weakness based metric M with Z . The function $M(Z(\bullet))$ is a metric on the space of protections.

The difficulty here is that mapping between protections and mitigated weaknesses exists in a theoretical sense, but in a practical sense the mapping may not be obvious to the owner of the system. The mapping Z may take a known protection system and map it into a set of unknown weaknesses. This means that the function Z is not computable. For if it were computable one would have a nice way of discovering unknown weaknesses. One would just devise a protection system, apply Z , and then see what new weaknesses are identified. So, he may indeed have a metric able to compare two different protections, but may not be able to say exactly what weaknesses the protections are mitigating, or even if one mitigates more weaknesses than the other.

Even if one did have a viable protection based metric M , what value could it provide? Suppose that $P_1 \subseteq P_2$ are two protections. It should follow that $M(P_1) \leq M(P_2)$. However, knowing that P_2 provides a higher level of protections than P_1 says nothing about how secure the system is. With both protections, the system may still have a catastrophic weakness that allows an adversary to exploit the system at will. The underlying issue here is identical to those surrounding consistent weakness based metrics. The metrics simply do not measure the things that one needs to measure in order to ask questions about security.

A *quality metric* measures how well a protection measure was implemented. These types of metrics can be applied with a certain amount of rigor. One can make comparisons across systems and check to make sure that the implementation meets certain standards. However, without other information about the system, these metrics cannot assure that any particular weakness has been mitigated. Since they are based on some standard process, they would have to be developed somewhat independent of the system and its adversaries. And as described above, it may not be able to map the protections to any particular weakness.

5.3. Threat Based Metrics

One method for designing a system with security in mind system is for the owner to start with his best guess as to his adversaries, their knowledge and resources. The owner simply chooses a resource bound high enough to exceed that of his perceived adversaries. Then the owner identifies his assets and what he wants to protect. The security system is built with those parameters in mind. In essence the owner estimates the set of exploitable weaknesses, E , and builds protections to ensure that the estimated value, E' , is the empty set. A *threat based metric* is a security metric whose inputs include estimated adversary capabilities and estimations on the set of exploitable weaknesses.

If one truly has a good notion of one's real adversaries, then this design approach makes sense in terms of practical security. However, as with weakness or protection based metrics, estimated measures of E do not give a true measure of how secure one's system is. One may severely misjudge one's real adversaries and their resources. Such a metric

may be definable and computable, but would be very brittle to changes in adversaries and their capabilities.

More importantly a threat based metric is based on the owner's *perception* of how well the owner has identified the security weaknesses and placed the proper protections to mitigate those weaknesses. Because it must be based on what the owner knows or believes, such a metric cannot measure how well the owner *actually* did their security job. The metric cannot, *a priori*, measure new threats that may arise.

On the practical side, a protection system may be much better at protecting the system than the system owner knows, by mitigating unknown weaknesses. This is often done by design. Security system designers often attempt a layered approach to security; they build something stronger or better than what they think they need, just in case. This overbuilding of the protection system often does mitigate weaknesses that would not otherwise be mitigated and, may in fact, mitigate unknown weaknesses. There are examples where it has been found later that if certain aspects of the protections were not in place, that the system would be vulnerable to newly discovered attack methodologies. By focusing on protections and making those as firm as possible, one may indeed make a secure system. One simply cannot measure that security.

6. Conclusions

There are two different ways to view axioms. One may view them as the definitions of the logical system under review. As such, they are correct and irrefutable. From them follows whatever follows. On the other hand, when attempting to create a logical construct that mimics reality, one then uses the axioms as a basis for comparison of the logical construct and the reality one is trying to mimic. In that case, the value of the resulting logic follows from the quality of the axioms. One must acquire some level of moral certitude as to how well those axioms match with reality.

We have proposed four axioms that seem reasonable and plausible and fit with the best-of-practice security reasoning available today. As with all axiomatic reasoning, one must choose to accept the axioms or not. If one accepts them, then it is clear that the notion and industry of defining all encompassing security metrics will have limited value at best.

On the other hand, if one is to reject the presented axioms, that rejection should come with a logical argument as to why they do not fit with reality. That argument, if valid, would be a core deviation from classical security thinking and would likely be quite valuable.

It is our belief that the axioms will stand. The implications are that it is not possible to measure the security or the trustworthiness of a real communication system in any sort of absolute sense. Efforts in defining trust metrics need to be focused in the practical arena. New definitions and notions can and should be developed that focus on the quality of what one has done, or to focus on the mapping between protections and the weaknesses that are mitigated.

In conclusion, one must realize that the practical side of security is still alive and well. After all, if, by luck or by design, no real adversary can exploit a system, that system is secure. The issue discussed in this paper is that even though a system may be strong, there is no way to measure its real strength and there is no way to know that your system is secure. Note also we do not claim that *all* security metrics must be trivial. There may indeed be great value in attempting to measure aspects of the system that can be identified. Application of the proper metrics may guide one's design efforts and assess the quality of an implementation etc.

7. References

1. INFOSEC Research Council, "National Scale INFOSEC Research Hard Problems List", November 30, 2005, http://www.infosec-research.org/docs_public/

DRAFT