All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

The harder problems

Since 1999, a “hard problem list” for information security has been published by the InfoSec Research Council. In 2005 it identified, in order: (1) Global-Scale Identity Management, (2) Insider Threat, (3) Availability of Time-Critical Systems, (4) Building Scalable Secure Systems, (5) Situational Understanding and Attack Attribution, (6) Information Provenance, (7) Security with Privacy, and (8) Enterprise-Level Security Metrics. Here's my problem...

When you ask the wrong question, the answer doesn't matter

These seem to me to be the wrong problems – mostly. I agree that (2) addressing insider threats is pretty hard, and that's one of the reasons I work on it. But (1) global scale identity management just isn't that hard. In 2005, it could have been done technically by using federation between nation states and creating a proper hierarchy. (3) Availability of time critical systems is a matter of redundancy and adequate resourcing. (4) Building scalable secure systems is predicated on the notion that we can already build unscalable secure systems, which we cannot. (5a) Situational understanding is a problem because we lack understanding of what we want to understand. (5b) “attack” attribution isn't the hard problem – the hard problem is “attribution”. (6) Information provenance is very much the same as the attribution problem. (7) Security with privacy is a misnomer because the terms are poorly defined – there is a tradeoff between attribution and privacy and between confidentiality and availability, etc. (8) “Enterprise level” can be eliminated without loss of meaning – we don't know how to measure”security” and we don't, on the whole, even know what security means.

What are some of the right questions?

Here's what I think is a better unstructured list: (old list elements in parens where applicable)

- Identifying and agreeing on terms for what we misname and overload as “security”.
- Identifying and addressing equities of different parties affected by “security”.
- Informing decision-makers of the implications of their decisions. (3, 5a, 7, 8)
- Attribution (of acts to the actors that undertake them). (5b expanded, 6)
- Identifying loyalties, extents, and changes in time to limit consequences. (2)

Some notions about the difference and a better direction

I intentionally avoided making a proper list to fulfill the real need because I thought some things might be clearer this way. Note that the InfoSec Research Council list is computer oriented, while my list is more people oriented. This is at the heart of the real “security” problems we face. Technology may create enormous benefits and bring with it enormous challenges, but that doesn't mean the way to meet those challenges is all about technology. Most of the challenges are from people and their behaviors, and most of the solutions are likely to be tied to people as well. I think we need to build a science and engineering discipline surrounding protection (keeping people from harm) and that it needs to address the human issues if it's going to be effective.

Building a protection science and engineering discipline

Fundamentals of science include the development of notions of causality. While we may want to have a non-causal existential reality, we don't really know how to do science without cause producing effect through mechanism. Assuming we will want the benefits of science, we will have to pay the price of assuming causality and seeking better understanding of mechanisms and how to use available cause to produce desired effect. That brings us to some start at identifying the challenges and clarifying them for now.

The challenges in more detail:

Identifying and agreeing on terms for what we misname and overload as "security"

I use the term "protection" (keeping from harm) as opposed to "security" (feeling safe) and more specifically, information (symbolic representations in the most general sense) protection.

There are many other terms that are increasingly used and have historically been defined. But there is a tendency to let folks slip on their usage and the result is sloppiness that produced poor science. Consider "hackers caused security leaks by phishing". Of course the first term could be anyone you want it to be, the second term is nonsensical unless whoever did whatever they did got unauthorized information about the protective scheme, and phishing really means lying. "Mandy lied to find out what was behind the door." is one interpretation, while another is "Chinese intelligence sent false email messages to NATO headquarters guards to gain information about their shift times." What is often meant is something more like unknown parties used an unidentified mechanism to cause unknown employees to reveal unauthorized information. But somehow, being more careful and precise about language is something we avoid in the information protection field, perhaps because we are concerned about leaking sensitive information, like that we don't really know what we are talking about.

In any case, we need to start using the already agreed upon definitions instead of making up words and misnaming things all the time if we are going to make progress as a field. And this indeed is a hard problem. The problem is one of educating a population, including the folks who think they already know something. It means people talking to and in the media using the words that are defined in the field, publications requiring the same words be used for the same things as a condition of publications, and yes – the word police will be out in force. It means making a profession out of a set of black arts. And that's a hard problem.

Identifying and addressing equities of different parties affected by "security"

The equities issue is often misunderstood or not addressed at all. It is another human issue. It has to do with the notion that what helps you may hurt me. In order to make decisions about protection, we are always trading benefits and harms. In many cases we are trading well identified benefits with potential and nebulous harms, and that is one of the big problems we face. The unrealized harm from an unexploited weakness is only a potential harm, while the lost opportunity from not doing something that could have resulted in that harm is very real.

For example, the World has gained enormous benefit from the Internet, but the Internet is inherently vulnerable to computer viruses because is increases sharing, is able to give others information received, and allows unlimited functionality. We have known since 1985 at least that this implies it is impossible to prevent all computer viruses. And we have rationally decided that the benefit outweighs the harm. The attackers like it, and so do the users. The defenders may not feel that way, but their equities are not well considered or understood.

Where equities really come up today is when one branch of government and society could benefit from capabilities of another branch, but cannot use them. For example, national intelligence agencies could have far better intelligence against actors in the US if they could surveil within the US as they do elsewhere. And of course law enforcement could likely do better if they had all of the capabilities of the Federal government behind them. But in order for law enforcement to have those capabilities, the sources and methods used might be revealed, resulting in the loss of the capability – perhaps as well as the lives of some sources. At the same time, law enforcement cannot share the details of their confidential sources with national intelligence agencies because it might be considered a violation of the constitution, and of course the possible loss of the lives of law enforcement informers.

I may think it's better to have a society in which governments cannot keep secrets, and you might think the opposite. But in the end, individual items of information are either limited or not limited in access. If the government has no secrets, and they pay for your health care in retirement, then we cannot meet both equities. The fees from your health care relate to conditions you have and the conditions can be derived from the fees. If the fee information is available, so is your health status. Claiming that we can control such things is problematic for 2 reasons; (1) the controls involve some level of secrecy and (2) limiting access while facilitating access produces an ability to derive the individual information from the aggregates. If we try to balance the two, we may lose both.

In any case, the equities issues are social in nature and deal with the nature of the different views of different parties regarding their desires and their missions. But they are also limited by our technical knowledge and ability to achieve the desired outcomes, the costs involved, and other similar things that play into the balance of the equities.

Informing decision-makers of the implications of their decisions

Since, in the large, nobody is able to perfectly predict the future, decision-makers make decisions based on limited information. Decisions almost always involve various acts in the present and future intended to cause and avoid various effects in the future. Notionally, better informed decision-makers make better decisions, even if there may be a definitional challenge or two in that notion. Today, we don't generally know what constitutes better, and that implies that we don't know what would be a better decision or what better informed might look like.

I personally favor more relevant and more certain information over less relevant and less certain information, but it's hard to tell what may end up relevant in hindsight, and certainty has never been easy to come by. In fact, we don't really even have a way to measure these things very well, and we don't really have scales designed for that purpose.

In an attempt to get at this, I have taken the approach that, in terms of information protection decisions, there are relatively few available alternatives for each of about 100 key decisions that often have to be made (or if they are not made, they are made de-facto). I started defining the decisions in terms of the alternatives, differentiated the alternatives in terms of the conditions that would cause each to be taken as opposed to the others, and build up a set of standard criteria for decision-making. This seems to work, in that it then drives the need for information. In particular, you can set about gathering the information necessary to make the decisions based on the decision criteria. Since the information is fairly specific and finite, it's not that hard to get it, and you can make decisions. But there remain a few small problems. You can make decisions, but the implications may remain largely unclear.

Understanding the implications of decisions is inherently a goal of risk management. That it, management of “risk” as opposed to “people”. Managing people involves working with them to get things done. It implies a feedback system in which people provide information that gets mixed with objectives by the manager who produces information for people to cause the system to meet the objectives. Risk management also implies a feedback system in which managers seek information about possible futures, mix them with objectives, and make decisions that will effect the actual future. But that's most of what is really understood. For events that meet the requirements of probability, there is a field of probabilistic risk analysis, but that is problematic when dealing with intentional actors, and it only deals with one facet of the problem space. A major challenge is that risk management is largely about anticipating and constraining behaviors of people and groups. This is a social science issue.

Attribution (of acts to the actors that undertake them)

If we could attribute all acts to the actors that undertake them, we could respond to all acts by holding the responsible actors accountable. This would have a deterrent effect on some and for others would mean that they get caught and punished for the things they do. If this is always true, then every bad act worthy of punishment or offensive to anyone has the potential of being punished, and from the day you are born till the day you die, every little mistake will be noticeable and if anyone is resources and interested enough to find it out, held against (or for) you. Note I didn't say anything about a court of law. Attribution in some sense means no more secrets. For everything that is done, those responsible will be clearly identified in their roles in it getting done. Rewards for your help in building the business, punishments for your imperfections.

The thing is, humans are terrible at this, and computers are even worse. We know some of the reasons people are bad at this. It's because we make all sorts of cognitive errors. This is the subject of study in psychology and sociology. Computers are even worse at this because they don't have the rich tapestry of inputs and experiences that people have to draw on. They simply accept information they are given and usually don't check anything out. They don't understand anything really, and they only normally fuse information from the sources directly used to provide it. For example, when I sit here typing, the computer doesn't watch my eyes, listen to my sounds, smell my fears, and so forth. Am I being coerced? Am I sleepy and not thinking straight? Am I having an argument with my boss? Am I even at the keyboard at all?

Identifying loyalties, extents, and changes in time to limit consequences

Loyalties to many different parties and ideas simultaneously exist within people. Nobody is probably 100% loyal to anything at the expense of everything else, and we likely wouldn't want people like that working for us. When people approaching this level of loyalty to, say a religion, are around, they tend to have what I would call poor judgement. They can be fooled into thinking that someone else is as loyal as they are and that the other person is more informed than they are, and as a result, end up following orders to commit suicide while blowing up a bunch of innocent (by my standards) bystanders. If I detected someone that was overly loyal to a single thing I would likely suspect them and try to avoid working with or being near them... unless of course they were loyal to me... just kidding. So the first problem we face is that we don't even know what we are looking for in loyalty. That's because, with a few notable exceptions, we don't study it as a field. One such exception is the US government, which studies those who have turned against the country when placed in positions of trust.

The work of PERSEREC is largely responsible for producing the adjudicative criteria used to determine who is eligible for security clearances. The criteria is based on the “whole person principal” in assessing eligibility, and it is imperfect. It is also based on studying the history of the subject and ruling out decision factors that were previously used, like sexual orientation, while identifying new correlating behaviors, like computer misuse. The basis is statistical studies of the past and analysis of the people involved in past cases. It is a social science.

The nature of this field is complicated, among other reasons, because there are relatively few identified cases that have been though enough of a legal process to be treated as reliable. In addition, there is the common challenge of understanding base rates and the inability to do any real experiments. The base rate problem stems from the fact that even if we figure out that 85% of those who illicitly reveal classified information are men, that doesn't mean we can use sex as an indicator. For example, what if 95% of the people with access to the information were men? Then women would have a disproportionately high rate of revealing it. But what if the total sample size was only 100 cases? Then the statistics would be so low that the differences in proportion would not be statistically significant. To study this subject properly, we would need to have far more disloyal people. But this might not be the right approach.

This issue of loyalty is highly subject to external conditions. People who would seemingly never steal anything would likely steal food if their families were starving and there was no other way to get it. Their loyalty to a set of beliefs, laws, etc. shrink in comparison to their loyalty to their children. That's why kidnapping family members works so well as a method to get people to do things they wouldn't otherwise do. In any case, this seems like an awfully hard problem and a really important one, given that insiders are responsible for more losses than any other identified group according to many studies over many years.

Still more hard problems?

In fact we have lots of other hard problems to address in the information protection arena. Here are a few that I haven't detailed above.

- **Why:** Computers don't have an ability to deal with the reasons for things. Acts that would be acceptable in some situations are not acceptable in other situations.
- **Any is not All:** We need people to be able to access anything, but not all things. Controlling how much of what things to access is something we don't know how to do.
- **Overhead:** All of the things we want to do better take time, effort, mindshare, and other overhead. We can only do so much before we refuse to take it anymore.
- **Dealing with imperfections:** Most protective mechanisms fail under a single fault. They are brittle to human imperfections and occasional bad judgement.

I'm not running out of hard problems. And I'm not likely to any time soon. But there is some commonality here, and it is worthy of repetition.

Conclusions

Most of the hard problems in information protection relate to people. We know how to do lots of things with computers, and in most cases, we can get computers to do what we instruct them to do. They are reliable and repeatable. But they are also inflexible and not well aligned to the humans that they serve. If we want them to serve us, we need to adapt them to our nature and not keep asking us humans to adapt to their nature.

