DEVELOPING THE THREAT MEASUREMENT HABIT

By Dr John Leach

This is the second in a series of three articles exploring a common theme, that simple security data can often be remarkably valuable from a risk management perspective. The challenge is to know how to use that data to get at the value it holds.

The first article explored the idea of creating incident profiles and showed how even rudimentary incident profiles can provide useful insights into where and how countermeasures are failing. This second article will explore the idea of measuring threats and will show how bringing the available data together and looking at it the right way can provide essential guidance to help security practitioners ensure their limited security resources are being applied where it will make the most difference rather than being scattered too thinly. A modest amount of threat data compiled through some relatively simple measurement steps is all that is required. Given how crucial it is to get a measure of the threats being faced, it is odd that there is not a lot more threat measurement being done.

FOLLOW THE DATA

The preceding article showed how to form incident profiles. An incident profile is the rate at which incidents occur not just in total number but as a function of some relevant parameter. Security incidents arise when the threat wins the contest between it and the countermeasures in its way. An incident profile can be interpreted as what is left after the threat profile has been pared back by the countermeasures in its way. The first article showed how, using this interpretation, even rudimentary incident profiles can offer up useful insights about where security countermeasures are failing. The ability to get useful insights from what might be no more than rudimentary data is especially fortuitous given that, often, people do not have a huge amount of robust incident data to work from.

If countermeasures are working well, there will always be a very much larger number of attacks than there are resultant incidents. Incidents will be relatively few and far between, and those incidents which do arise will be the exceptions, the ones where the particular attack, for whatever reason, just happened to defeat the countermeasures in place at the time. Hence, incident data is almost certainly going to be more sparse than threat data and cannot always be relied upon to give a representative picture of the threats which were their cause. Incident profiles will always be based on fewer data points than the threat profiles from which they came.

This suggests that, perhaps as an alternative approach, it is to the threat profiles one should turn. If that is where the bulk of the data is, maybe threat profiles can be used to get a different angle on the dynamics leading to risk and maybe they can provide different insights complementing the messages provided by incident profiles.

John Leach has been an Information Security professional for more than 20 years. He has held senior positions in the security teams of a number of organisations, and in December 2002 formed his own consultancy company to allow him to pursue R&D in his specialist area of risk modelling.

John Leach has an academic scientific training and long experience working with national and international organisations in the security field. He has brought these together to create innovative solutions to complex problems and to develop a unique range of consultancy services around the theme of using security data scientifically to advance our understanding of how security risk is created and to quantify how the countermeasures we apply reduce the risks we face.

John Leach has been an active member of the Management Committee for the Information Assurance Advisory Council (www.iaac.org.uk) continuously since May 2002. He is also a member of the International Board of Referees for *Computers and Security*.

THREATS PROVIDE ESSENTIAL CONTEXT

Besides there being potentially more data available if measuring threats rather than incidents, there is a second, quite compelling reason for wanting to measure threats. Everyone will be familiar with the adage "you can't manage well what you don't measure!". In a similar sense, one can't expect to protect well against threats which are not measured. Security is about protecting valuable assets from the threats that bear down on them. Surely it is essential for those doing the protecting to have some understanding of the size and nature of the very threats they are trying to resist? Nobody wants to spend a huge effort protecting their assets from threats that are more figment than fact, and everyone needs to be sure they are protecting their assets against the threats which are real and substantial. In other protective domains, safety for example, threat measurement is second nature.

It is somewhat remarkable that the security industry has never really developed the threat measurement habit. This can be illustrated quite simply by looking at security surveys. There are now several well known security surveys conducted regularly, usually annually or biennially. They provide all sorts of statistics about the number of organisations in each size band that suffer this or that common security problem. And they present, alongside those figures, further statistics about the number of organisations that do or do not practice each of a number of basic security practices (such as having a security policy or updating their AV regularly). But, puzzlingly, these surveys contain almost no information at all about the threats causing the common security problems on which the surveys are based.

Security surveys such as these are the equivalent to, say, safety surveys talking about people who fall down stairs but which don't say what it was about the stairs which led to these people falling. These surveys are providing the security equivalent of how many people suffered just minor bruising from their fall and how many people were off work for two weeks with a broken ankle, and what these and other people do to ensure they don't fall and hurt themselves more often, but are not providing any explanation about what caused these people to fall in the first place. Were their stairs rotten through? Were their stairs uneven? Slippery? Did they have spillages or obstructions on them? Were they simply too steep? It seems clear enough that the measures needed to prevent accidents on stairs must reflect the reasons someone might fall in the first place. There is little point going to the expenses of redesigning the safety features of a staircase if the main problem causing people to fall is that they run up the stairs speed reading the e-mails on their PDA rather than taking the time to walk carefully.

As this illustration shows, threat measurements provide the context justifying which countermeasures are needed and which are little more than wasted effort. By measuring what it is that makes a threat more likely to lead to accidents or incidents in one place than in another, threat measurements can show what it is that makes a countermeasure effective. Threat measurements enable security practitioners to ensure they are focussing their limited security budgets on doing the things that will make a real difference rather than having to scatter their resources thinly in a vain attempt to cover all the bases.

DESCRIBING THE THREAT / COUNTERMEASURE CONTEST

If measuring threats is essential to providing context, then why has the security industry, unlike other protective areas, not yet developed the threat measurement habit? There are possibly two reasons. The first is that it is not always clear what to measure. The second is that there is a widespread belief that measuring threats requires enormous infrastructure and effort, beyond the resources of most ordinary organisations. Both of these barriers can be removed.

The first of these barriers was addressed in the previous article in this series, from the point of view of incidents. That article talked about how to organise incident data in a way which brings out the incident profile, the way in which the rate of occurrence of incidents rises or falls as a function of one or more relevant variables. It then described how to work out which were the variables needed to form each incident profile. A very similar line of reasoning is used for measuring threats and creating threat profiles.

To summarise, for each threat of interest identify the key countermeasure used to protect against that threat. For this key countermeasure, work out what it is that makes that countermeasure either more or less effective at protecting against that threat. That will show the "contest" taking place between the threat and the countermeasure and then take you to the variable which controls the effectiveness of each attack in that contest.

For example, if the threat of interest is e-mail viruses and the main countermeasure is AV software on the mail gateway, then the primary variable which makes that AV software either more or less effective against viruses is the frequency with which the software checks for new virus signatures. The "contest" between the AV countermeasure and the e-mail virus threat is the race between a new virus making its way to the mail server and the AV software updating itself with that new virus' signature. The previous article showed how this view of the contest led to identifying the key threat variable, the main variable governing whether or not the threat won the contest with the countermeasure. For this example, it was the age of the virus (measured in hours) at the time the mail server is first exposed to it.

This line of reasoning can be used to identify the key threat variable for any threat / countermeasure contest. This is the variable that needs to be used as the basis for measuring the threat. Some further examples are given below.

THE FIRST RUNGS OF THE LADDER

Identifying what it is about each threat which needs to be measured clears the first of the two barriers to developing the threat measurement habit. The second barrier, that somehow threat measurement requires enormous infrastructure and effort, is an understandable miscomprehension. The examples below serve to show that it is possible to get useful results with just a relatively small amount of effort. Naturally, the more effort put in to threat measurement, the more detailed and comprehensive will be the insights obtained from it, but there is certainly no reason why anyone should feel it is necessary they build a fully comprehensive hugely detailed measurement infrastructure. One doesn't need to have one's own Hubble telescope to enjoy looking up at the stars. Indeed, very very few people will ever need to reach nearly so far. Early X-rays and ultrasound scans were useful diagnostic tools even though the images they produced were very blurry. Simple pictures can provide a good indication of where risk is being created even if they aren't highly precise.

Think instead in terms of progressing up the rungs of a ladder. The first rungs are easy to climb, and each step takes you that little bit further and allows you to reach that little bit more. And how often when you get a ladder out at home do you ever need to climb all the way to the top? You go as far up the ladder as you need to go and no more. This analogy holds for measuring threats.

The first rung of the threat measurement ladder involves gathering from amongst the already available data and takes little more effort than that needed to get started. Already available data can be generic data available from external sources about the threats at large, or it might be data already being created as a by-product from existing countermeasures or operations. The next rung might involve one or two simple steps to filter out from the generic data that data that pertains to your particular circumstances or to show how the threat varies from office to office or system to system. The third step maybe does involve a little dedicated machinery to create data that otherwise wouldn't be created. But you won't climb up to that rung unless you judge it is within reach and you believe it is needed to get you up to where you want to go. Each step up the ladder takes only a small amount of additional effort and brings you a corresponding amount of additional insight. Where you find that insight is of great value, you will refine your measurements and advance on to the next rung. If it is unclear what the current rung's results are telling you, you will hold steady at that level until you have decided what to go for next. Start with simple and modest plans, grow them no more than you need them to grow, and you should expect to get a return on your efforts with each step you take up the ladder.

SOME SIMPLE THREAT MEASUREMENT EXAMPLES

Having established that threat measurements provide essential context, and that the data ladder allows anyone to work towards the results they want in a manageable way, this article will now provide some examples to show how the approach can work.

The first step when measuring threats is to identify one's primary security objectives. Assuming your security goal is to provide appropriate protection for key business assets, then what might be the main security objectives underpinning your goal? They might include some of the following:

- Ensuring Business Continuity
- Preventing Unauthorised Physical Access to Restricted Areas
- Preventing Unauthorised Logical Access to Systems or Data
- Preventing Unauthorised Disclosure of Sensitive Proprietary Information
- Maintaining the Security of your voice systems

For each security objective, choose the main threat of interest and the main countermeasure used to protect against that threat. Thinking in terms of the contest between the threat and the countermeasure will indicate how the threat should be measured given that countermeasure. For example, for the above security objectives, the threats and countermeasures might be:

Objective	Top Ranked Threat	Main Countermeasure
Ensuring business continuity	Major environmental disasters (e.g. hurricanes damaging production facilities) and international crises (e.g. international borders closed to transport in response to avian flu pandemic)	Business Continuity Planning
Preventing unauthorised physical access	Actual or suspected unauthorised entry to a restricted area which could mean the integrity of a key operational process has been compromised.	Physical access controls (e.g., locks on windows and doors, positive identification of all personnel entering a site, monitoring all entry and egress into and out of key restricted areas)
Preventing unauthorised logical access	Actual or suspected unauthorised access to critical systems or sensitive data which could lead to system or data compromise.	User account access rights reviews. To a prescribed frequency (e.g. every six months for high-risk systems), all the user accounts which have access to a given system have their logical access rights reviewed and inappropriate or unnecessary access rights removed.
Preventing unauthorised disclosure	Attempts to steal valuable proprietary information, either maliciously or opportunistically, enabled by the information's not being adequately protected according to its sensitivity.	Information classification and labelling
Protecting voice systems	Toll fraud and other misuse or abuse of the company telephone system	Voice Security Audits

Now, thinking in terms of the contest between each threat and its countermeasure, identify the variable which determines how much of a danger each attack within the threat stream might be. Then look at how to measure the threat in terms of that variable. Take each of the above security objectives in turn.

Ensuring Business Continuity

To understand the risk from this threat it needs to be measured, for each type of disaster or crisis (and the main types of disaster or crisis which need to be considered might vary with location), in terms of the rate of occurrence of events as a function of the magnitude (usually duration) of the event.

A number of public and private sources provide (free or for a subscription) a threat rating for each country or location. This reflects where the threats they cover occur most often. Choose the threat rating scheme which most closely reflects your main concerns and the source which covers the countries or locations where you have major sites.

This will give you an initial threat profile. Augment this with regional and local statistics from industry associations, commercial and public sources relating to the type of incidents which most often arise (physical, logical, accidental, environmental), their frequency and magnitude, and their financial or other impact, for each geography where you can find the data. This will let you start to flesh out the threat profile. Where data is available, correlate the impact statistics to the size of the operation affected and the level of company or site preparedness.

This threat measurement can tell you which of your sites need to have the most robust BC plans and profiles the major types of event which realistically happen so you can test local BC plans sufficiently.

Preventing Unauthorised Physical Access to Restricted Areas

To understand the risk from this threat it needs to be measured, for your business' critical protected processes, in terms of the rate of occurrence of attempts to breach an individual physical access control, whether successful or not, as a function of the effort (or time) it would take for an adversary to breach that control.

Rate each of the main types of physical access control you use according to the level of effort required by an adversary to defeat that access control. This can be a broad strength rating system, rating controls on a scale of, say, 1 to 5. You should expect to find that those at the lower end of the rating scheme get tested more often by staff than those higher up.

For each main type of physical access control, measure the number per month of attempts (failed or successful) to breach the physical control. The opportunity to measure this will vary with the control and will need to be balanced with the desire to keep the impact on the organisation to a minimum. If locked doors are monitored, measure the number of access card/PIN combinations which are rejected (filter out instances of simple keying error), and reconcile the number of staff who enter a monitored area against the number of staff who leave it. Count the number of personal access passes lost (whether accidental or stolen) and the number of times physical access controls are found to have been compromised (e.g. doors not being properly locked, staff seen entering in groups on a single person's access pass, unaccompanied visitors found in restricted areas) regardless of whether the compromise of the control is suspected to have been malicious or careless, and regardless of whether the compromise led to any harm.

This threat measurement can tell you whether you have a significant problem with unauthorised physical access and, if you do, can tell you whether you should put your effort into getting staff to change behaviours or into strengthening windows doors and locks.

Preventing Unauthorised Logical Access to Systems or Data

To understand the risk from this threat it needs to be measured in terms of the rate at which staff attempt unauthorised logical access as a function of the value of the system (high Integrity systems) or data (high Confidentiality data) at which the improper access is aimed.

The threat comprises those attempts to access systems for which the user should not have the requisite access privileges. For some of these attempts, the user will (wrongly) have the requisite access rights and for some (correctly) they will not. Those where the user wrongly has the requisite access rights will not appear as failed access attempts in the relevant logs. Hence, it is necessary for all access attempts to be logged, successful and failed, not just failed access attempts.

For each main operating environment, from the relevant logs, measure and set a norm for the rate of failed access attempts (filter out failures which were obviously caused by factors other than knowing attempts by users to gain unauthorised access, e.g. system misconfigurations). Then, as each system has its user accounts' logical access rights reviewed and redundant or inappropriate rights removed, measure how the rate of failed access attempts varies in the months following the review. Whenever the review process identifies (and removes) a user account's excess access rights to a significant system and the user should have known their access rights were inappropriate, review the past three months' logs to see if the user knowingly used those inappropriate rights. This is the most direct way to measure this threat.

If you wish to climb one step higher, categorise each reviewed system (e.g. product development; production; sales; finance; HR; administrative) and then correlate the threat as measured above with the category of system, to understand which types of system are most targeted by the threat.

This threat measurement can tell you whether you have a significant problem with unauthorised access, and the nature of the systems and data being targeted the most. If you do have this problem, put significant extra effort into getting staff to change their behaviours. It will also indicate whether your process for granting logical access rights is robust or weak.

Preventing Unauthorised Disclosure of Sensitive Proprietary Information

To understand the risk from this threat it needs to be measured in terms of the rate at which conduct by staff handling sensitive information causes that information to be disclosed improperly, for each level of sensitivity of the information.

Unauthorised disclosure is not easy to detect making the rate of these attacks difficult to measure directly. Instead measure:

- General threat indicators. Encourage people to report examples of suspicious behaviour they experience (e.g. unexpected telephone calls asking for information not normally in the public domain, official-looking but unsolicited e-mails asking staff to complete a questionnaire about one or other aspect of the company's business) regardless of whether any information was disclosed.
- Specific threat indicators. Use scanning utilities on e-mail servers to measure the rate at which classified information is being sent in e-mails (internal or external) plus the rate at which classified information is being sent to questionable destinations (e.g. hotmail / yahoo accounts, partner / non-partner / competitor domains). This will establish the current level of each specific threat channel and help to indicate the general level of the threat.

To put these threat indicators into context, measure your vulnerability to unauthorised disclosure. Encourage people to report valuable information assets they come across exposed (e.g. sensitive reports left out in unoccupied offices, sensitive information about the company they come across on discussion lists) and use that as an indication of your level of vulnerability. Conduct office walkabouts after hours to identify sensitive information assets left exposed to disclosure.

Where possible, gather equivalent data for industry at large (from industry associations, commercial or public sources).

This threat measurement can tell you whether you have a significant problem with unauthorised disclosure of sensitive information. If you do not have an information classification scheme in

place, this will provide you with the information you need to make the business case for establishing a scheme.

Maintaining the Security of your voice systems

To understand the risk from this threat it needs to be measured in terms of the rate of attacks (attempts to exploit telecom system weaknesses) by type (internal/external; fraud/abuse/misuse) and as a function of the potential impact (e.g. increased call costs) caused.

Use industry figures (e.g. from the CFCA – www.cfca.org) to establish a norm for the rate and range of telecom system attacks (internal and external) that companies of your size can expect to experience. Toll fraud (higher than appropriate call costs) is the most commonly cited result of attacks but there are many other outcomes which can be substantially more damaging.

As you conduct your regular voice security audits, use the results to indicate which weaknesses tend to be most common in your voice systems. Use that information to guide your review of your Call Detail Reports (CDRs) and history logs (recording the administration and maintenance activity conducted) to identify the level of misuse, abuse or fraud which has occurred on your PBXs in the recent past. Track the changing profile of calls and the fall in your monthly call costs as PBX security is improved.

Industry figures indicate that, on average, a company saves around 5% of their monthly call costs with a straightforward tightening up of their voice system security. The more your monthly call costs fall as you tighten security, the more of a problem you have had in the past from this threat, and the more of a problem you are probably still having.

This threat measure can give you a very good indication of the magnitude of the problem you are suffering and how to address it directly. It is very important to get this problem under control before taking on Voice over IP and IP telephony generally, as the opportunities for adversaries to exploit voice system weaknesses multiply when moving from traditional voice systems to IP-based ones.

CONCLUSION

Though the security industry does not yet focus on measuring threats, there are clear reasons why threat measurement is effort very well spent. Threat measurements can provide essential context to enable limited security resources to be directed to where they can have the most effect, and even modest amounts of measurement effort can provide useful threat insights which would otherwise remain out of reach. You do not need to make major investment in measurement infrastructure and you might well decide that the first few rungs of the ladder provide a view perfectly sufficient for your needs. The examples given here lead to simple threat profiles, not detailed ones, and initial insights into the size and nature of the threat, not fully precise insights, but the results are nonetheless useful as a starter. Naturally, the more effort you put in to threat measurement, the more you should expect to get out in the form of detailed and comprehensive insights. But there is no law requiring you to go all the way. Start small and grow your measurement capabilities no more than you decide you want them to grow. You will get value for money at each step along the way.

As security is evolving to become less an art and more a science, as it undoubtedly is, threat measurement will shortly become the norm. The sooner you get started, the sooner you will see the benefits.