UNCOVERING THE VALUE WITHIN SECURITY INCIDENT DATA

By Dr John Leach

This is the first in a series of three articles exploring a common theme, that auditors often have access to security data which can be extremely valuable from a risk management perspective. The challenge is to know how to extract that value. In general, security professionals are not practised at marshalling the data resources they have to hand and at interpreting data to uncover the insights contained within. So much value is sitting just below the surface but most people do not have the scientific training to enable them to put their hands into that data and pull out the gems buried within it. These three articles will give three separate views on how to do just that.

In this issue, the first article will explore the idea of creating profiles. It describes how to use incident data to create incident profiles, and how even rudimentary incident profiles can give useful insights into where and how countermeasures are failing. The second article will look at what might be called "the data ladder". It is a common misapprehension that one needs to climb to the top of the ladder, expending a lot of effort building a large data gathering operation, before anything useful will come out of the data gathering exercise. It turns out that climbing even the lower rungs of ladder can give useful insights in return for just a modest amount of effort. Finally, the third article will look at how to use existing compliance results to show much more than just the level of compliance. Without having to redo the assessments, it is possible to show system and business owners the nature of the risks their systems are running and stimulate a constructive discussion about risk reduction rather than just have an argument around the need to improve compliance.

INTRODUCTION

Whether you already have a growing repository of incident data or are planning to start a new incident reporting programme, incident data can be a mine of critically useful risk intelligence and insights. The question is how best to tap in to that mine to extract the wealth of risk insights buried within.

There is an inescapable feature of security incidents which, once it is understood, can provide a lead on how to set about extracting those insights. In this article, we will explore what causes incidents to be unique, and use that to get at the risk intelligence buried within security incident data and to build an incident monitoring programme which can make a real contribution to risk management and risk reduction.

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*.

THE UNIQUENESS OF SECURITY INCIDENTS

When you dig down into a repository of incident data, you quickly come to realise that each incident you come across is essentially unique. Different incidents relate to different attacks, different systems being targeted, different vulnerabilities left unclosed, and to different countermeasures failing in different ways. This uniqueness, along with the huge diversity of incidents

which can be experienced, are what make it so difficult to get to the root causes of what is going wrong and to the risks you are running beneath the surface.

This uniqueness of each individual incident isn't just nature's revenge, nature's way of making every worthwhile objective an uphill struggle. It is an inescapable feature of security incidents. Incidents are like people in that no two will ever be exactly the same. However, despite the uniqueness of each individual person, we have, by studying the people we see, learned a lot about the underlying forces which move and shape us. Similarly, we can learn a lot about the underlying risks we are running if we can understand how to go about studying the security incidents we see.

THE RELATIONSHIP BETWEEN INCIDENTS AND THREATS

It is an unavoidable consequence of this uniqueness of incidents that the value locked up in an incident repository is not to be released by drilling down into the detail behind each individual incident. It is to be released instead by reading what incidents can tell us collectively. Incidents need to be studied in aggregate, not just in isolation.

This leads us immediately to our first question: given the need to aggregate our data, how should we do that in order to analyse it? For this, we need to understand the dynamics governing the relationship between incidents and threats.

A threat can be thought of as a stream of multiple attacks grouped together, despite their variability, by the common theme or aspect they share. That common aspect may be broad and general or it may be narrow and specific, depending on the particular threat being studied at the time.

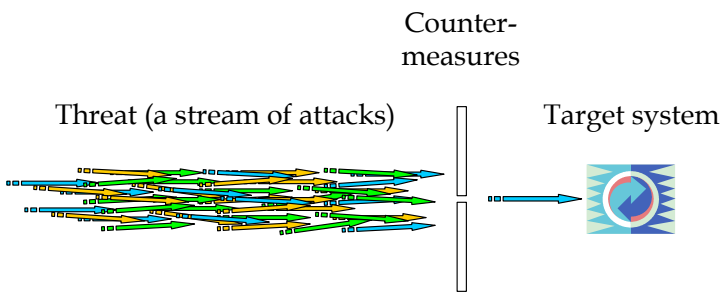

Figure 1 - Incidents arise from the interplay between a threat and the target system along with its countermeasures

Each attack in the threat stream is often, for the most part, independent of any other attack. The variability from one attack to another comes about because the specific nature and timing of each attack is influenced by a number of factors some of which cannot be known or predicted. For example, nobody can predict the paths a newly released virus will follow as it spreads around the globe, or when the opportunity will next present itself for a staff member to commit a fraud. It is these unknowable and unpredictable factors which make some of the details of each attack we see random and governed by chance.

It is this random variability which leads to each incident being unique. Each time an attack within the threat stream just so happens to win out, to get past the countermeasures in place at the target system, an incident arises. And just as each attack within the stream will be slightly different from any other attack, the precise details of how each successful attack defeats the countermeasures and then engages with the target will vary. Random variability in the attacks within a threat leads to random variability in the incidents caused by the threat.

AGGREGATE BY THREAT

Though it might not be immediately obvious, this does give us the lead we need for aggregating incident data. Often, if one has a collection of data which displays some form of variability, a useful ploy is to organise that data in a way which leads to some form of profiling. This is what we should do here.

Incidents are the form in which the outcomes of successful attacks manifest themselves, and attacks are the manifestations of the threat they represent. Hence, each incident is representative of the threat which gave rise to it even though, alone, each incident is no more than a very poor indicator of that threat. It is not until you look at the incidents collectively that they become less a collection of random incidents and more a representative indication of the range of possible outcomes caused by the threat.

This tells us that incidents need to be aggregated according to the threat which causes them. Doing this will allow us to use incident data to learn about the magnitude and nature of the threat and about the way the threat defeats the countermeasures which stand in its way. From there, we can get an understanding of the underlying risk posed by that threat and how to influence or manage that risk.

The first step, then, is to aggregate incident data by the threats of interest. Because often there might be only a limited volume of incident data available, it is usually best to start this from the perspective of broadly defined threats. The analysis can always be progressed to work with more narrowly defined threats later if the quality and quantity of incident data will allow.

At the initial, broad level, the threats to consider would be threats such as:

- Staff knowingly violating a (written or unwritten) code of authorised behaviour (e.g. e-mailing a confidential file to an unauthorised external party; committing fraud);
- Staff innocently performing security-relevant errors of commission or omission (e.g. accidentally overburdening a system, causing it temporarily to lose availability; inadvertently setting off a fire alarm which leads to an evacuation);
- External untargeted IT attacks (e.g. viruses and worms);
- External targeted attacks against one's customers (e.g. phishing);
- External targeted IT attacks against one's systems (e.g. attacks against web applications);
- Denial of Service attacks (e.g. traffic flooding);
- Natural accidents or failures (power surges; power outages; hardware failures);
- Environmental threats to persons or property (e.g. earthquake, fire, flood, lightening);
- Man-made threats to persons or property (e.g. arson, bombing, extortion, kidnap);
- And so forth.

Aggregate the available incident data according to the relevant broad threats. This will allow incidents to be aggregated into broad groups where, within each group, we would hope to have enough data to perform some statistically meaningful analysis. For example, we might group together all incidents which arise from staff knowingly violating a code of authorised behaviour. This would include incidents of staff fiddling their expenses through to incidents of major fraud, staff making overly expensive private telephone calls on company time through to staff using the company-provided desktop to pull down and distribute illegal content from the web. Large or small, these incidents are all manifestations of the same broad threat, the threat of staff knowingly breaking the rules. The exact form in which staff break the rules is immaterial at this initial stage. The extent to which they break the rules is not. We'll return to this point in a moment.

PRACTICAL LIMITATIONS IMPOSED BY THE DATA

Before we talk further about how to use this aggregated incident data, we should take a brief diversion to talk about the practical constraints and freedoms the available data allows. If you have been able to gather a large amount of incident data, you will not be limited to doing your analysis only at this broad level of aggregation. If you have a limited amount of data, the level of detail you can get down to, and the degree of certainty of the risk insights you produce, will necessarily be correspondingly limited.

The level of detail to which you can go, increasingly partitioning your incident data into smaller groups corresponding to more narrowly specified threats is determined by two considerations: how much data you have available and the particular risk management problems you are trying to address. It is probably safe to assume that everyone will want to take their analysis down to as detailed a level as the available data will allow. Ultimately, the limit comes when there is insufficient data in each group to allow meaningful analysis. How much data is needed depends on the quality of the available data and that is often hard to measure. Therefore, as a general rule of thumb, start with the broad level of aggregation described above. When that analysis has been completed, try taking the analysis down a level, breaking your broad threats down into their constituent parts and partitioning your data into constituent subgroups. Your data will determine how far down you can go before the results start to become too uncertain to be of real use.

Of course, these levels of aggregation and analysis are not mutually exclusive. You can do both the broad analysis and the more granular analysis in parallel on the same repository of data, the different levels will just provide different insights.

THEN, PROFILE BY COUNTERMEASURE

The way to use the aggregated incident data to get at the available risk insights is largely independent of the level of granularity you are working with. Hence, what follows here applies equally to both broadly and narrowly defined threats.

When you look through the incident data in each of the groups you have formed, you will see that the data spans quite a wide range of incidents. The next question this brings us to is this: how do we arrange the data in each group to expose the particular insights it provides? The answer is that we need to organise the data in a way which brings out the incident profile, the way in the number of incidents recorded varies as a function of one or more relevant parameters. The key is to work out which parameter or parameters to use to generate that profile.

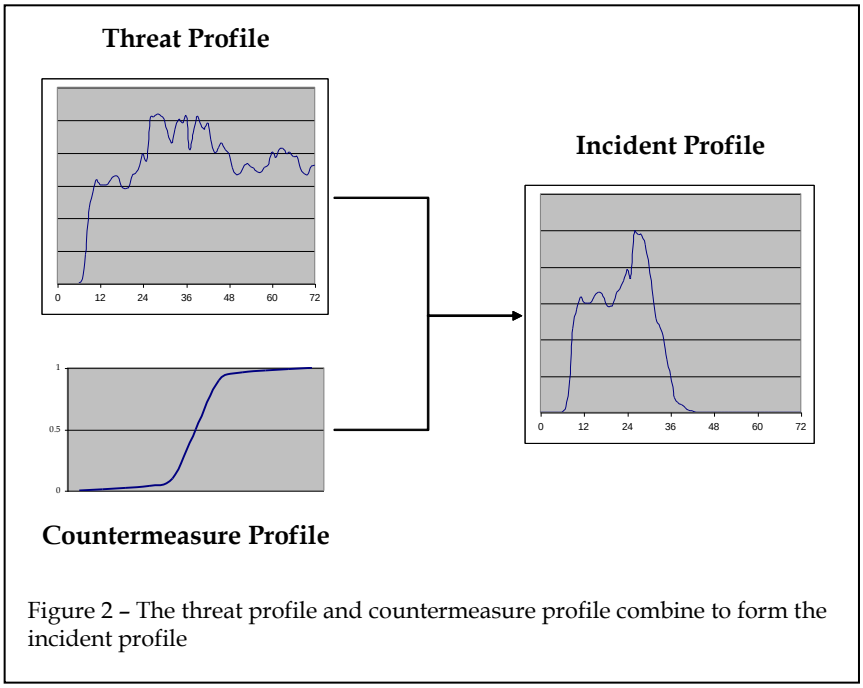

Figure 2 – The threat profile and countermeasure profile combine to form the incident profile

As shown schematically in Figure 2, an incident profile can be thought of as the combination of two other profiles, a threat profile and a countermeasure profile. The threat profile is the number of attacks within the threat stream as a function of a relevant threat parameter. The countermeasure profile is the effectiveness of the countermeasure (between 0 and 1, or between 0% and 100%) as a function of the same threat parameter. Any countermeasure is going to be more effective at stopping some threats than it is at stopping others, and usually a little bit of thought about the countermeasure will tell us what that relevant parameter is. At this stage, it is not necessary to be able to calculate how effective the countermeasure is as a function of this threat parameter, just to be able to work out what the relevant threat parameter is.

One at a time, for each threat of interest, identify and rank the key countermeasures you use to protect yourself against that threat. There might be several countermeasures you use so try to

identify the one which has the most significant effect. Start your incident data analysis with this one countermeasure in mind. (At a later stage, assuming there is sufficient incident data available, the analysis can be refined to bring a second or third countermeasure into the analysis.)

For this one countermeasure, work out what it is that makes this countermeasure either more or less effective at protecting against that threat. This is often easier for technical countermeasures than non-technical ones, so start with technical countermeasures to develop some practice.

Let's look at an example to see how this might work. If the threat of interest is e-mail viruses and you decide that your main countermeasure against that threat is the AV software on your mail gateway, then you might consider that 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 more frequently it checks, the more likely it is to have a given virus' signature in its local signature file and thereby be able to recognise and block that virus when it sees it. The less frequently the software checks for signatures, the more likely it is that the virus will not be recognised when it comes in on an e-mail and the virus will be let through by the AV software. The contest between your AV countermeasure and the e-mail virus threat then becomes a race between a new virus making its way to you and your AV software's updating of its signature list. If the new virus can get to you while it is still very young, before your AV vendor has released a signature and your software has had a chance to download that signature, the virus will win. If your software can download the virus' signature before you get exposed to the virus for the first time, you will win. In this example, then, the threat parameter of interest, the main parameter governing the threat / countermeasure contest, is the age of the virus (measured in hours) at the time you are first exposed to it.

Identifying this key parameter allows you to construct an incident profile from your incident data. Interpreting the shape of that profile is how you uncover the risk insights buried within.

THE INTERPRETATION OF PROFILES

An incident profile is created by ordering the incident data according to the relevant threat parameter. Staying with the example above, you would order your e-mail virus incident data according to the age, at the time you were exposed to it, of the virus which caused the incident. There are various sources you can turn to for when a new virus was first recorded, and your incident logs will tell you when the virus incident occurred.

The result might be an incident profile similar to the one in Figure 2 above. This shows us that the number of virus incidents is very small for very young viruses. We can assume this is because there is only a very small probability of anyone being exposed to a virus when the virus is very young. Looking at how the incident profile changes moving along the axis towards older viruses, the number of incidents grows up until the point where the AV defences start to take effect. Thereafter the number of virus incidents falls. The number of incidents due to very old viruses is very small, this time because most AV defences can be presumed to be nearly 100% effective against old viruses.

A 1-dimension incident profile of this form can provide some directly useful risk management insights.

- If the profile is highly peaked (as in Figure 2), that would indicate that a small improvement in that countermeasure's effectiveness could lead to a large reduction in the number of incidents seen. If this is the type of incident profile you see, your risk management priority should be to review your standards for deploying that countermeasure to see if there is a way you can improve its current effectiveness.
- A profile which peaks and then has a pronounced tail would indicate there are some places within the organisation where that countermeasure is operating much less effectively than it operates in other parts. This would suggest that security standards are not being applied uniformly or that the countermeasure has not been deployed fully across the organisation. If this is the type of profile you see, look at where within your organisation or infrastructure the tail incidents are coming from. This will tell

you where the countermeasure is being least effective. A comparison between the deployment standards followed in those locations and the standards followed elsewhere will allow you to work on getting all parts of the organisation up to the level of effectiveness of the best. For some countermeasures, usually the less technical ones, it is not easy to find a good way to measure the countermeasure's effectiveness. Being able to use the tail of an incident profile to do that in this way could be the key.

- A generally broad incident profile with no obvious peak would indicate that this countermeasure is always going to be of limited effectiveness against that threat. If that is what your incident data shows, your risk management priority should be to focus on building defence in depth and strengthening other countermeasures, in particular strengthening your mitigative countermeasures (those which reduce the severity of the disruptions caused by each incident, e.g. having a speedy incident response capability) or your alleviative countermeasures (those which reduce the degree of harm an incident with a given severity can be expected to cause, e.g. having contingency plans in place, or insurance).

This discussion shows that teasing an incident profile out of incident data can immediately provide some clear risk management insights, insights which might otherwise have remained buried away. The source incident data does not need to be hugely accurate or complete for these purposes, at this stage it is the overall shape of the profile which is important.

THE INTERPRETATION OF PROFILES (II)

Of course, when interpreting incident profiles, one must also consider what other influences might be at work on the data. One factor which is important for manually reported data is unevenness in reporting. Some staff will be more inclined to report incidents than others. For example, staff working with hazardous materials or in a bank branch environment might be more inclined to report incidents than, say, staff working within an IT environment. Some staff might report some types of incident but not others, thinking for whatever reason that some incidents are of less importance for reporting than others.

Two ways to tackle reporting inconsistencies come to mind. Firstly, provide reporting guidelines and share reporting practices around the organisation so that all staff can understand what is expected of them and see what others are doing. Another way is to use data which is believed to be reliable as a check on other data which is thought to be less reliable. This technique can be very powerful if used well. A broad threat, such as staff knowingly violating a code of authorised behaviour, is made up from a number of sub-threats which you can reasonably expect will follow similar overall lines to each other. If you can get a reliable incident profile for one of the sub-threats, then you can use that to set expectations as to the general shape of the incident profile you should be seeing for each of the other related sub-threats. If the sub-threat incident profiles do not follow the same general shape as each other and there is no obvious reason why this should be the case, then you might presume you have a data quality problem with some of your data. If so, the more reliable incident profiles can be used to adjust for quality problems in the less reliable profiles, allowing the less reliable incident data to make more sense.

Another factor which hampers the interpretation of incident data is data incompleteness. You might have be given access to largely historical data, have decided which threat parameter is the most relevant, and then found that the incident data provided to you doesn't contain the relevant data fields. This is as frustrating as it is commonplace, and it crops up most often with incident data which was gathered without anyone having thought through what in particular it was important to record for each incident. If you starting to plan your incident gathering programme or are about to relaunch one that has not worked well in the past, then take the opportunity to think through beforehand how you might want to analyse your incident data. It will prove enormously valuable in helping make sure you capture all the relevant data for each incident.

And, thirdly, bear in mind that threats might vary across the organisation or with time. Technology threats can be volatile, changing with time as the technologies to which they relate change. Staff

related threats can vary significantly with culture and hence, for international companies, with office locality. Incidents gathered over a long period of time or for a wide variety of locations will reflect these influences and your incident data analysis will need to take these factors into consideration.

MOVING A PROGRAMME FORWARDS

How far and how deep the incident data analysis can be taken depends largely on the amount of data available. The value of the risk management insights which can be formed when reliable incident profiles can be produced, profiles reliable enough to allow meaningful assessment and interpretation, justifies starting with broad threat definitions and not rushing into the detail. If the threat definitions used are too specific and each aggregation produces a sparse scatter-diagram rather than a defined incident profile, then the data is being overstretched.

If there is plenty of incident data available and good quality incident profiles have been produced, only then should you try taking the analysis on to more granular threat definitions or to 2-dimensional incident profiles in order to release a deeper level of risk insight. If the data is not there to support that level of analysis, then you should focus on building up both the quality and quantity of data over time.

While putting that building programme in place, there is another course of action you can pursue in parallel to supplement the incident data you have, and that is to measure and profile the threats themselves. In some situations, it is easier to build a reliable threat profile than it is to build a reliable incident profile because there can be a lot more data available with which to build.

This will be the subject of the next in this sequence of three articles.