Intrusion Detection Systems as Evidence

Peter Sommer
Computer Security Research Centre,
London School of Economics & Political Science
P.M.Sommer@lse.ac.uk
Phones: +44(171)955 6197; +44(181)340 4139

Peter Sommer *is Senior Research Fellow at the Computer Security Research Centre, London School of Economics & Political Science where his research speciality is Legal Reliability in Information Systems and where he teaches Information Systems Security from a social science and public policy perspective. His commercial consultancy includes expert witness work in both criminal and civil proceedings, computer forensics, insurance risk analysis and investigations, "best practice" in document management and e-commerce. He is currently Specialist Advisor on Electronic Commerce to the UK House of Commons Select Committee on Trade and Industry.*

This paper was originally presented at the RAID 98 Conference, Louvain-la-Neuve, Belgium
<http://www.zurich.ibm.com/~dac/Prog_RAID98/Full_Papers/Sommer_text.pdf> *and in an updated version to the NIPC Conference, New Orleans, Sepetmber 1999*

Intrusion Detection Systems as Evidence

Peter Sommer
Computer Security Research Centre,
London School of Economics & Political Science
P.M.Sommer@lse.ac.uk

ABSTRACT: Although the main aim of IDSs is to detect intrusions to prompt evasive measures, a further aim can be to supply evidence in criminal and civil legal proceedings. However the features that make a ID product good at providing early warning may render it less useful as an evidence-acquisition tool. An explanation is provided of admissibility and weight, the two determinants in the legal acceptability of evidence. The problems the courts have in dealing with novel scientific evidence and the differences between "scientific" and "legal" proof are discussed. Criteria for the evaluation of IDSs as sources of legal evidence are proposed, including preservation of evidence, continuity of evidence and transparency of forensic method. It is suggested that the key to successful prosecution of complex intrusions is the finding of multiple independent streams of evidence which corroborate one another. The USAF Rome Labs intrusion of early 1994 is used as a case-study to show how defence experts and lawyers can undermine investigators' evidence.

One desirable post-event outcome of a successful detection of an intrusion may be legal proceedings. This was the observation of the NSTAC Network Group Intrusion Detection Subgroup [1] in December 1997:

> An additional concern the subgroup identified was developing intrusion detection capabilities that are responsive to the needs of the law enforcement community. Collecting information and protecting the chain of evidence of computer intrusions so that it will stand up in court has always been a challenge for system administrators and law enforcement. ….
>
> Specifically, the subgroup found that:
> - Current intrusion detection systems are not designed to collect and protect the integrity of the type of information required to conduct law enforcement investigations.
> - There is a lack of guidance to employees as to how to respond to intrusions and capture the information required to conduct a law enforcement investigation. The subgroup discussed the need to develop guidelines and training materials for end users that will make them aware of what

This paper was originally presented at the RAID 98 Conference, Louvain-la-Neuve, Belgium <http://www.zurich.ibm.com/~dac/Prog_RAID98/Full_Papers/Sommer_text.pdf> *and in an updated version to the NIPC Conference, New Orleans, Sepetmber 1999*

information law enforcement requires and what procedures they use to collect evidence on an intrusion.

I want to address these problems: what is required to turn the output of an Intrusion Detection System (IDS) into legally reliable evidence?

At the outset we need to recognise that there are significant and valid aims of IDSs which have nothing to do with legal proceedings. At the same time the ability in legal proceedings to demonstrate that a specific intrusion has occurred may be merely one step in proving guilt of a criminal offence or liability in a civil matter. The problems lie in the detail, the nature of the gap between the purposes of IDSs of various types and the needs of the legal system. The gap exists not only at a purposive and functional level but also in philosophic approach, particularly in relation to what the computer community on the one hand and the legal system on the other think constitutes "proof". In practice rather more is needed than "guidelines and training manuals".

To map the geography of the gap we need a bit of background: we need to revisit the various types of IDS and examine the sort of outputs they produce; we need to look at what happens in legal proceedings and at general notions about the nature of evidence, particularly "scientific evidence"; then we can attempt to identify the points at which bridges may usefully be built across the gap.

Types of IDSs and their outputs

The main aim of IDSs is to detect intrusions. Depending on the precise IDS, typical hoped-for outcomes can include:

- the ability to react in a timely fashion to prevent substantive damage – by automatic or manual intervention
- the ability to react in a timely fashion to mitigate substantive damage – by automatic or manual intervention
- the ability to identify activity which is the precursor of a more serious attack
- the ability to identify a perpetrator
- the ability to discover new attack patterns
- indirectly to provide an additional measure of system protection beyond that available from other forms of security measure
- and of course, evidence

In many respects these sometimes limited aims for systems which detect intrusions into computers echo those for devices which detect attempts at physical intrusions into buildings: in the domestic and low-end commercial/industrial sector, devices which respond to the breaking of an infra-red path, or changes in temperature or pressure, or to movement, are useful in that they can be used to trigger an alarm which deters the intruder; however they have no immediate means of collecting any evidence which might identify a specific intruder so that they could eventually be charged. A simple closed circuit television recording system which is not constantly monitored may collect visual evidence of visitors to a particular location – evidence which is potentially useful in identifying individuals – but by itself it has little value in providing a timely warning of intrusion. Most premises security depends on a combination of methods.

If we now look at the main classes of IDS we can see their limitations more clearly. The table below follows a fairly conventional taxonomy which divides IDSs into systems which operate after the event and rely on the analysis of logs, audit trails and the like, and those that attempt

real-time monitoring in the hope that unusual patterns of activity give sufficient reliable indication that evasive action is possible before real damage occurs. This real-time monitoring can either be sited at the computer system which is the putative target or placed on a network where traffic can be evaluated. The taxonomy also distinguishes between IDSs that react to a known body of signatures of supposedly intrusive behaviour and those that use heuristics and AI to identify anomalous activity. The second column summarises the apparent advantages and disadvantages of each type of IDS:

Method	Advantages / Disadvantages
Post-event Audit Trailing – from library of signatures	no possibility of immediate reaction but likely to produce useful anomaly detection which could lead to further action; depends on quality of library of signatures; however potentially useless if audit trail compromised
Post-event Audit Trailing – by detecting unusual patterns: statistical anomaly detection	no possibility of immediate reaction but depending on volume of data and sophistication of eg AI tools likely to produce useful anomaly detection, but also likely to produce false positives; potentially useless if audit trail compromised
Real-time monitoring of packets on the network link – against library of signatures	permits real-time alarm and thus evasive action; depends on quality of monitoring tool and library of signatures, volume of data and location of monitoring point(s)
Real-time monitoring of packets on the network link – by detecting unusual patterns: statistical anomaly detection	permits real-time alarm and thus evasive action; depends on quality of monitoring tool, volume of data and location of monitoring point(s); runs risk of false positives etc
Real-time monitoring of activity on host /network device- against library of signatures	permits real-time alarm and thus evasive action; but alarm may not sound early enough; depends on quality of monitoring tool and library of signatures – otherwise potential for false positives and negatives; problems of data volumes, compromise of overall system performance; danger that host is compromised and with it monitoring tools / audit logs
Real-time monitoring of activity on host /network device – by detecting unusual patterns: statistical anomaly detection	permits real-time alarm and thus evasive action; but alarm may not sound early enough; depends on quality of monitoring tool – otherwise potential for false positives and negatives; problems of data volumes, compromise of overall system performance; danger that host is compromised and with it monitoring tools / audit logs; runs risk of false positives etc

The advantages and disadvantages reflect the following criteria (which have a degree of inter-dependence):

- Timeliness of warning given
- Extent of freedom from false positives / false negatives
- Completeness of coverage
- Availability of signatures (intrusions, methods, ways of recognising them in terms which can be recognised by tools) or reliability of anomaly-detection methods
- Speed and quantity of traffic to be monitored vs speed of IDS tool
- Scalability
- Ease of configuration
- Ease of use
- Fault tolerance
- Resistance to compromise / subversion
- Cost of product
- Cost of additional computer resource / impact on overall system performance

- Cost of administration
- Relevance/Contribution to overall security policy / planning

Of course a number of current ID products seek to combine several methods.

But from an evidential point-of-view what we seek is something we can demonstrate to others long after the event itself is over: that tends to mean logs of various kinds. Potentially these can include:

- system logs
- audit logs
- application logs
- network management logs
- network traffic capture
- contemporaneous manual entries

In addition it is also possible to process this primary data into a form in which it is easier to analyse and understand. We can call this "derived data". Indeed there are products, originally designed to aid in cases of narcotics trafficking and serious fraud, which offer highly visual presentations of complex data so that patterns of activity become apparent.[2]

Even before we get to the specific hurdles erected by the needs of the legal system these logs may be deficient in terms of their ability to persuade a third party:

- the logs may make little immediate sense without training in the operation of the IDS tool and an understanding of the principles upon which it operates
- the logs may lack sufficient detail
- the logs may not exist over a sufficient time period for comparisons of normal and abnormal activity to be made
- the logs may be incomplete for the relevant period of time
- in the case of real-time monitoring the monitoring tool may not be able to keep up with the stream of traffic with which it is expected to deal
- in the case of real-time network monitoring the network location of the device hosting the monitoring tool may be such that it is unable to capture all relevant traffic, some of the packets using other routes
- the logs may not sufficiently distinguish between a legitimate and an unwanted access
- the logs may not identify the perpetrator in any useful way
- the logs may have been compromised *prior* to collection as potential evidence
- the logs may have been compromised *during* collection as potential evidence
- the logs may have been compromised during *post-collection* analysis
- in the case of derived data, the methods of analysis and subsequent presentation may lead to misleading results

Evidence in Court

We now need to change direction to understand a few things about the nature of evidence as it comprehended and dealt with by the courts. Evidence doesn't exist as an absolute by itself; it is material which is used to establish the truth of a particular fact or state of affairs.

These remarks refer to those legal systems broadly follow the English common law model. Whereas in Continental Europe the criminal procedure sees investigations being carried out by a specialist judge – *juge d'instruction* – in countries like England, the US, Australia and many former members of the old British Empire, investigations are carried out by the police or other law enforcement agency, the decision to prosecute is made by a separate body – District Attorney in the US, Crown Prosecution Service in England, and at trial the role of the judge is as chairman of the proceedings and enunciator of law. Separate opposing legal teams represent the arguments of prosecution (the Crown, the People) and the defence. The trier of fact is a jury. *

The procedure, known as adversarial, has lead to the development of complex rules of evidence, describing what can and cannot be put before the court for its consideration of fact. In common law countries, "proof" is what courts decide to accept as true after due consideration of evidence presented before it. The standard of proof for private, civil disputes is the "balance of probabilities" and for criminal prosecutions "beyond a reasonable doubt". In law evidence is no more or less than that which tends to persuade the court to a particular conclusion. With a few exceptions, the demonstration of proof has to take place each time and for each individual set of circumstances.

Evidence[3] has to satisfy two tests: **admissibility** (that is, it must conform to certain legal rules which are applied by a judge) and **weigh**t (that is, it must be understood by, and be sufficiently convincing to the court - whether there is a jury or a judge acting as a trier of fact).

The rules of **admissibility** have their basis in the history of the jury system and the distinction that must be made between witnesses and jurors – those who have seen something and those specifically appointed to decide what has happened. One of the best known rules is "hearsay" – Bob tells Alice Charley has confessed to committing murder – all Alice can tell a jury is that Bob has reported to her a conversation he has had with Charley; Alice has no evidence herself that Charley has confessed. Another traditional rule says that a document is not admissible by itself but must be produced by a live witness who was responsible for it; only then can the court take note of its content. This particular rule is substantially eroded; even in the 19th Century it was recognised that in the case of bank statements it would be impractical to call all the clerks that might have made an entry in a particular ledger and in England, for example, we now admit "business documents and records" without too much burden. [4]A further typical rule excludes opinion evidence unless from an identified expert. The "best evidence" rule, also being eroded, favours the original over the copy – a matter of some difficulty in relation to documents held in or created by a computer. In criminal law evidence may be excluded on the grounds that it has been acquired unfairly or illegally; in the US the basis is the Fourth Amendment; in England exclusion is a matter of judicial discretion. [5] In the US most of the law is found in the Federal Rules of Evidence and in England in the Police and Criminal Evidence Act, 1984 and Civil Evidence Act, 1995. But in both jurisdictions there are also many cases which provide detailed interpretation. One writer on the law of evidence said: "Founded apparently on the propositions that all jurymen are deaf to reason, that all witnesses are presumptively liars and that all documents are presumptive forgeries, it has been added to, subtracted from and tinkered with for two centuries until it has become less of a structure than a pile of builders' rubble." [6]

In non common law jurisdictions – as in Continental Europe and all those countries formerly colonised by Continental European countries – the principle is largely of the free admissibility of evidence. But then the criminal procedure is not adversarial but investigatory.

* This is not true of all criminal procedures but a clear distinction is always made between a finding of fact and an interpretation of law

Once evidence has been admitted, the court can then assess it for **weight**. Juries are not under any obligation to follow any specific method nor, with a few very technical exceptions, are there any categories of evidence which they are compelled to accept.

The admissibility/weight distinction sometimes becomes quite unclear and particularly so in relation to "novel scientific evidence". Most of the evidence that comes from modern computers is likely to fall into that category. In England the assessment of this is largely a matter of weight and cases are lost if an expert and a judge appear to be telling a jury that they *must* accept a finding. Thus in R v Doheny, [7] DNA statistical evidence was produced in a rape and buggery case. The conviction was overturned an appeal on the basis that the expert had overstepped his role. "He should properly, on the basis of empirical statistical data, give the jury the random occurrence ratio, the frequency with which the matching DNA characteristics were likely to be found in the population at large…. That would often be the limit of the evidence which he could properly and usefully have given. It would then be for the jury to decide, having regard to all the relevant evidence, whether they were sure that it was the defendant who had left the crime stain, or whether it was possible that was left by someone else with the same matching DNA characteristic."[8]

In most US states however, it is the judge who acts as a gate-keeper in evaluating whether the expert evidence comes from "generally accepted scientific principles". This was established in the classic 1923 case of Frye v US [9] and re-evaluated and extended in Daubert v Merrell Dow [10] in 1993. *Frye* has been criticised because in any one case it may be unclear what generally accepted scientific principle is at issue, because it assumes that there is a coherent and identifiable scientific community and that it inevitably reaches the "truth". *Daubert*, a case which involved a drug that caused birth defects was an interpretation of Federal Rule of Evidence 702: "If scientific, technical, or other specialized knowledge will assist the trier of fact to understand the evidence or to determine a fact in issue, a witness qualified as an expert by knowledge, skill, experience, training, or education may testify thereto in the form of the opinion or otherwise." *Daubert* produced four tests: (1) whether the theory or technique can be (and has been) tested; (2) the error rate associated with the method; (3) publication in a peer-reviewed journal; and (4) whether the technique has gained widespread acceptance. And again it has been subject to extensive criticism.[11]

Once a *Daubert* test has been passed the actual evidence can be safely presented to a jury. The difference in approach between the English and US Federal jurisdictions is partly attributable to the greater use of juries in the US compared to the UK. Apart from defamation cases, juries have all but disappeared from UK civil trials and in criminal matters only the more serious cases are heard by magistrates sitting alone.

The notion of legal "proof" is thus only distantly related to that of "scientific proof". While "scientific" proof depends on the application of generally recognised methods of scientific investigation, "legal" proof depends more on the rules of admissibility of evidence and what is convincingly presented in court.

In fact there are many types of evidence that may be brought before a court:

- **real**, that is, an object which can be brought to court and examined on the spot
- **testimonial**, the eyewitness observations of someone who was present and whose recollections can be tested before the court
- **documentary**, a business record in any form which once proved as to authenticity can be examined for content
- **expert**, the opinions of someone expert in a particular field and/or the conclusions of that expert after carrying out a specific investigation

- **derived**, a chart, video, etc., created from primary evidence to illustrate how certain conclusions might be drawn

The logging types of evidence that IDS produce directly are documentary, but they require the testimony of people who were responsible for setting up the logging systems and for collecting the specific logs which the court is being asked to consider. In addition an expert may be required to provide context, explanation and interpretation. And it may help if charts etc are produced

While in the final analysis "weight" is a non-scientific concept, nevertheless there are a number of desirable features that one looks for in non-testimonial evidence – that is, exhibits and documents of various kinds. These attributes include that an exhibit is:

- **authentic** – that is, specifically linked to the circumstances and persons alleged
- **accurate** - free from any reasonable doubt about the quality of procedures used to collect the material, analyse the material if that is appropriate and necessary and finally to introduce it into court - and produced by someone who can explain what has been done.
- **complete** - tells within its own terms a complete story of particular set of circumstances or events

In relation to more technical types of evidence – forensic evidence we can expand on the range of attributes:

- there should be a clear **chain of custody** or **continuity of evidence**; for more normal types of evidence found at a scene of crime this means that the circumstances in which the item was found are clearly described and perhaps photographed *in situ* and that there is a complete account of all that has happened to the item up to the point at which it is being presented in court. This should include a search record at the scene of crime, an entry in a search register held at, eg a police station, any specific handling by law enforcement officers or forensic scientists. Many law enforcement agencies use a bag-and-tag protocol – the item is placed in a bag using a one-time numbered sealing tag; each time the bag is opened the sealing tag is changed and all involved keep separate records in their daily books
- a forensic method needs to be **transparent**, that is, freely testable by a third party expert. This creates difficulties if law enforcement think disclosure might result in the design of counter-measures which would prevent its future use; again the devisor of a forensic procedure may find it difficult to maintain commercial confidentiality
- in the case of material derived from sources with which most people are not familiar quite extensive **explanations** may be needed
- in the case of exhibits which themselves contain statements - a letter, database record or other document produced by a computer , for example - "accuracy" must encompass the **accuracy of the process** which produced the statement as well as **accuracy of content**; normally that requires the document's originator to be make a Witness Statement and be available for cross-examination

Another feature of the procedure is that the defence can demand disclosure (sometimes called discovery) of material which the prosecution may have used by which they have decided not to rely on in making their case.

Evidence in Practice

Most real-life proceedings are unlikely to be just about unwanted intrusions into computer systems.

This holds true even if we look at those offences which appear to be solely about "hacking". Thus in the English version of this sort of offence, s 1 of the Computer Misuse Act, 1990, before a conviction can be obtained, prosecutors need to demonstrate:

- that a computer was involved
- that it was accessed
- that such accessing was unauthorised
- that the person who is charged with the unauthorised access knew at the time that the access was unauthorised

The contribution an IDS can make to this prosecution only covers the first two items; indeed on some readings the IDS can only address the second bullet point.

But in more serious offences, the central feature may be the commissioning of a fraud, the stealing of information or an act of terrorism. In civil disputes, the fact of unauthorised access may be grounds for a suit of breach of contract or the tort of negligence, or it may be grounds for dismissal from a job (and the inability of the employer to prove the unauthorised access by an employee grounds for a suit for wrongful dismissal).

The classic requirement in the typical hacking case is to prove that the person charged is the perpetrator; in the USAF Rome Labs attacks of March and April 1994[12], the first offender used the handle of Datastream Cowboy and later took on the identity of more than one USAF officer. But the person in the dock at Bow Street Court in London was called Richard Pryce - how do you make the linkage and do so to criminal standards of proof: beyond a reasonable doubt?

The answer is that single streams of evidence (perhaps one logging file from one IDS) are unlikely by themselves to make the case; an individual stream may be vulnerable in that it falls foul of an admissibility rule or it may lack sufficient weight (because, for example, the log is not readily understandable or can't be independently tested or is otherwise deficient in crucial detail). The single stream may need to be corroborated by an separate independent stream of evidence which tells much the same story (the second witness of the same event). But this particular event may only be part of the sequence which must be proved before a conviction is secured (or breach of contract or of duty of care established). So we may need a multiplicity of streams of evidence, some from IDSs, some from other computer and telecommunications sources, some from human witnesses, and all hopefully corroborating one another.

In the case of Richard Pryce, his main *modus operandi* consisted of using a bluebox program on his PC (hard-disk capacity, 170MB) to abuse a freephone directory facility in Bogota, Colombia. From there he dialled into a Seattle-based ISP which offered free shell accounts. Here he called himself Datastream Cowboy and used the Unix work-space to run a collection of widely available hacking programs. His main target was the .MIL domain and for a while, through determination and stamina rather than any unusual skill he was very successful.

The streams of computer-derived evidence included:

- applications, text files, chat sessions with other hackers, lists of IP addresses, lists of cracked password files etc found on his PC
- a phone log from his telco showing numbers dialled, time and duration of each call
- logs of a USAF monitoring tool located in a shell account of one of the ISPs where Richard also had a shell account which he was using as a principal launch pad for his activities as "Datastream Cowboy"
- logs from a USAF network monitoring tool which captured IP traffic on a local USAF sub-net upon which some of the target computers were located
- evidence from some of the target computers
- a derived exhibit produced by Scotland Yard's Computer Crime Unit to show how all these exhibits could be linked together with clock-timing data to reach only one conclusion.

Many of the primary streams of evidence contained clock day/time data and this provided the means to lock the separate streams together to tell a complete story.

The slides show, respectively and in a very simplified form, the path of the hack by which "Richard Pryce" in London becomes the "Datastream Cowboy" who attacks USAF computer facilities, the monitoring tools; and the types of evidence that were finally available:

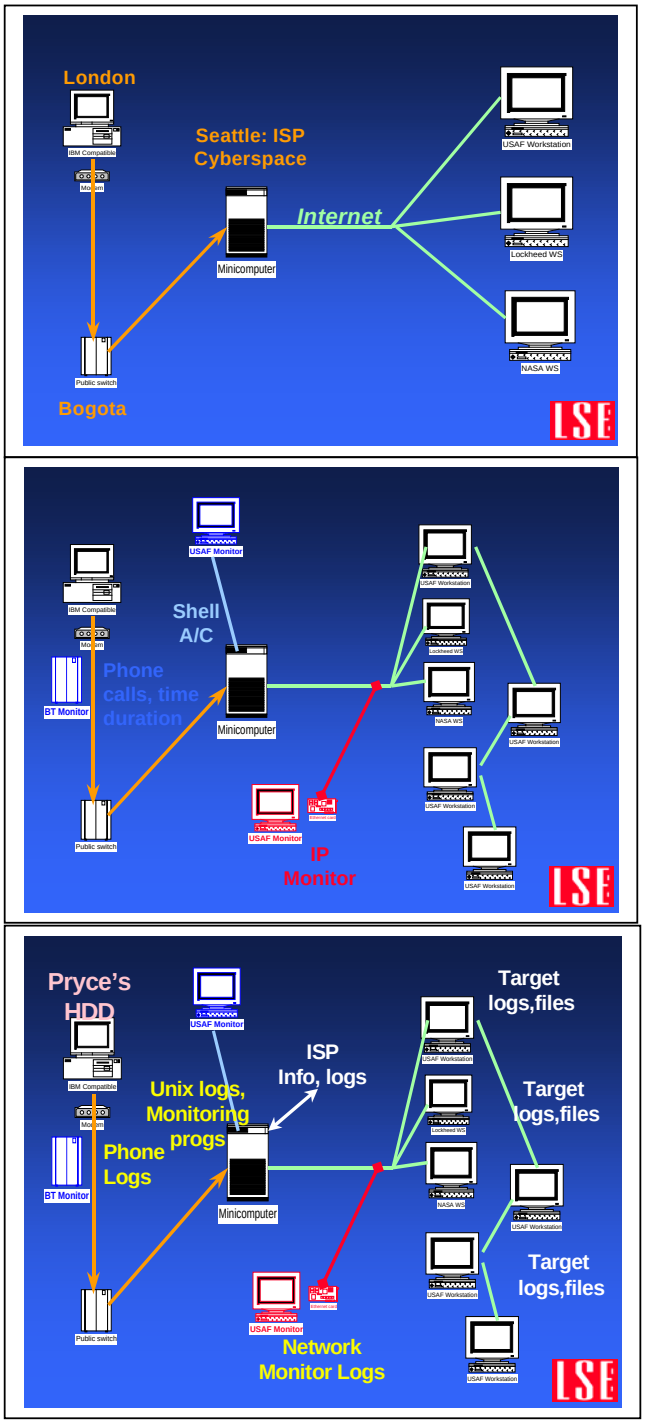
London
IBM Compatible
Modem
Seattle: ISP
Cyberspace
Internet
Minicomputer
USAF Workstation
Lockheed WS
NASA WS
Public switch
Bogota
LSE
USAF Monitor
Shell
A/C
Phone
calls, time
duration
BT Monitor
IP
Monitor
Pryce's
HDD
Target
logs,files
ISP
Info, logs
Unix logs,
Monitoring
progs
Phone
Logs
Network
Monitor Logs

In many ways this was in form a model investigation, though it was not intruder detection tools which gave the initial identification of Pryce; several months before the events for which he was charged, Pryce had spoken to a USAF intelligence officer on an IRC channel and given away some personal details including his home phone number. After being raided Pryce made admissions of a general nature about hacking. But very many months later, when the actual charges became known and turned out to be much more severe than he and his family thought appropriate he decided that he may have to fight. Still later and after the English approximation to plea-bargaining, Pryce pleaded guilty to a restricted range of charges so that the evidence was never properly tested.*

But during the period while he and his lawyers thought that a full trial might take place I was hired to test the quality of the technical evidence. In the English criminal justice system the defence can instruct its own experts; these have a dual role: initially to advise defence lawyers on the technical evidence so that they understand the issues and in turn can advise their client; secondly to give evidence in court. In this latter role (which does not follow automatically from the first) the expert becomes a special sort of witness but the duty is no longer to any client but to the court; questions whether of fact or opinion must be answered truthfully and completely; an expert cannot be simply a "gun for hire". The starting point for a defence expert is the list of specific charges which have been made and it is the evidence for these that is to be tested; the expert does not have a duty to correct the deficiencies in a prosecution case nor to form a view of the defendant's guilt.• It is unlikely that I would ever have given evidence *for* Pryce and my role was limited to testing the evidence.

And it turned out that, good though the structure of the investigation was, almost every individual stream of digital evidence could be challenged. For example:

- **material on the hard-disk of Pryce's computer** Following usual practice a complete sector-by-sector copy was made by non-invasive means and the copy was later analysed for incriminating material including some on deleted files that were recovered[13]. The most significant were printed out. But the usual practice also gives defence experts access to the sector-by-sector copy so that the forensic procedure can be tested and also to allow the defence to see if there is material on the disk which contradicts the view the prosecution wishes the court to accept. However in this case the prosecution claimed that Pryce had downloaded three files of considerable security significance; the sector-by-sector copy couldn't be released as that would have resulted, among other things, in the further release of the sensitive files. But the defence would have argued that the methods used to produce the print-outs of the files upon which the prosecution were relying could not be scrutinised, that this was prejudicial, and that as a result all the evidence derived from Pryce's hard-disk should be excluded.
- **logs of phone activity from Pryce's home** These logs were obtained from a device attached to Pryce's phone; they did not capture the content of traffic (difficult under the UK's Interception of Communications Act, 1985). The logs were said to show that unusual combinations of numbers were being transmitted and that some calls lasted for several hours. Taken together with the existence of a phone phreaking program on Pryce's computer, this was said to demonstrate

* The true identity of Pryce's co-accused, who called himself Kuji, was not established until a long time after Pryce was charged. The evidence against Kuji was always less complete and immediate and partly because of mistakes made in charging him, the case against Kuji, real name Mathew Bevan, was abandoned so that here again the digital evidence did not face substantive challenge.

• In fact experts *cannot* be asked the so-called "ultimate issue" question of a defendant's guilt.

that phone phreaking was taking place. However the call monitoring device appeared, from the print-outs supplied, to be inconsistent in it ability to record the unusual combinations of numbers: how far therefore could the print-out be relied on?

- **Activity at the Bogota Telephone Exchange** This was the next stage in Pryce's journey, but no evidence was adduced from there. The next point at which a trace becomes visible was at a Seattle-based Internet Service Provider called Cyberspace at which "Datastream Cowboy" had a shell account
- **Evidence from Cyberspace ISP** USAF officers also acquired shell accounts and they were able to monitor activities on Cyberspace with a tool that used combinations of *ps, who* and related commands. There were several problems with the logs produced in evidence: many of them contained log-in "welcome" pages from Cyberspace apologising for a recent hard-disk crash – how far, therefore, could we rely on the data that the Cyberspace box was now reporting? More seriously, the various exhibits (individual items of evidence) showed a break in the chain of custody: the print-outs were of e-mails containing the original logs, not the original logs themselves. The dates of the emails showed a significant amount of time had elapsed since the original events. There was no description of any anti-tamper controls. Worse still, the person producing the exhibits to the court was not the person who had collected the original exhibit (USAF had had to assemble quite a large team) - this was a clear violation of the hearsay rule and would have rendered all the logs inadmissible. Finally, initial requests by the defence for access to the monitoring tools to assess their reliability were refused; the defence would have claimed to be prejudiced by this.
- **Network Monitoring Tools** These, also installed by USAF officers, worked in the usual way: the ethernet devices on one or more Unix boxes were put into promiscuous mode and set to capture packets based on data within the header. The criteria were set on various combinations of origination and destination information. Elsewhere use was made of TRACERT and its variants. The problems with the resultant print-outs produced in evidence were very similar to the other USAF monitoring evidence. USAF did not want to release the source code for the tools so that it could be assessed for correct working; it is unlikely that they would have been happy to disclose the topography of the local sub-net upon which the monitoring tools and the target computers were located, though without that map it would be difficult to assess the tools for completeness of coverage. Here too the monitoring was a team effort but the witness statements did not properly reflect this and there were concerns about the lack of a fully accounted for chain of custody
- **Evidence from the target computers** These computers had been attacked, passwords, files and processes compromised. How then could we rely on evidence from these computers? No attempt had been made to create a proper trap in which a dummy "environment" was created as a honey pot to attract the hacker while other processes, carefully separated and guarded, recorded what was going on. Instead we were left with statements in which legitimate owners of computers "recollected" and "recognised" certain events and passwords.

Who can say how these arguments might have played in court? Interestingly enough one of the perceived problems of bringing to justice hackers who operate on a global stage didn't especially arise: there was excellent co-operation between US and UK officials and witness statements produced by US citizens were sufficiently carefully written as to meet some of the

eccentricities of the way in which English law requires computer-derived evidence to be treated[*].

The point of this brief prosecution assassination exercise (the print-out exhibits exceeded 10,000 pages and there was also a great deal of electronic evidence) is to show how quickly the output of a well-thought-out investigation can become vulnerable when subject to hostile scrutiny in the adversarial atmosphere of legal proceedings.

Re-designing IDSs as sources of Evidence

There are a number of useful conclusions that can be offered:

1. The value of an IDS depends, in the first instance, on the extent to which it provides timely and accurate information of the likelihood of an intrusion so that evasive action can be taken. On the whole, the earlier the warning is given the less likely it will provide details of the identity of the perpetrator. In terms of the security aims of most organisations, prevention or evasion of attack is preferable to post-event legal remedy or assisting law enforcement.
2. The carrying out of an investigation which leads to the identification of a perpetrator need not automatically result in the production of evidence that is admissible and believable by a court.
3. Evidence acquisition is a separate but related exercise. It is best carried out against a checklist which identifies the main problems of admissibility and where the main focus of the gatherer is court explanation and presentation.[14]
4. Single streams of evidence are unlikely to be adequate to convince a court; what is required are multiple independent streams of evidence which corroborate each other.
5. The feature that will link together these independent streams will usually be day-time clock data; some means of synchronisation is thus necessary.
6. If logs are to be produced from IDS tools, a prosecutor must be prepared to disclose complete details of the tool, and how it was configured and operated. In the case of network monitoring tools, the disclosure may have to include details of the topography of the local net.
7. Logging evidence, along with anything else that has been generated by computer, will need to be formally "produced" to court by the people intimately involved. Most jurisdictions have provision in their rules of evidence for the production of the results of team work, but these need to be followed in detail.
8. In the case of evidence logging which is to take place on a putative target, arrangements need to be made to prevent compromise during attack. It may be possible to do this within an OS on a single machine, but an alternative arrangement would be to send, as soon as possible after each event is recorded, all logging information to a separate, cryptographically-protected security vault[15].
9. Logging evidence should always be "best", that is, straight from the computer upon which it was created. Even if the logs are subsequently processed in order to perform an analysis or make them easier to understand, the raw log should always be available.
10. Where an exhibit is built from "derived" data (as in a chart or spreadsheet) the raw data has to be available for disclosure to the defence.
11. There needs to be a complete "chain of custody" or "continuity of evidence" from source to court. This can be done by statement, entries in note-books and registers, and also by using computer technology to make tampering more difficult: typically this can include

[*] For example, certificates of "proper working" under s 69, Police and Criminal Evidence Act, 1984

the writing of logs to WORM media at an early stage (this process itself should be properly recorded) or by using digital finger-printing.

12. Almost certainly the IDS market will be strengthened by the arrival of a range of procedures and products which concentrate on evidence collection and preservation, as opposed to intrusion.

---

[1] President's National Security Telecommunications Advisory Committee (NTSAC) Network Group Intrusion Detection Subgroup, available on http://www.ntsac.org as FIDSGREP.PDF

[2] For example the Analyst's Notebook from I2 Limited

[3] For a more extended legal perspective, seen particularly in terms of English criminal law, see Sommer, P, *Digital Footprints: Assessing Computer Evidence*, Criminal Law Review Special Edition: Crime, Criminal Justice and the Internet (1998) pp 61-78

[4] ss 23-24 Criminal Justice Act, 1988, s 8 Civil Evidence Act 1995

[5] s 78 Police and Criminal Evidence Act 1984

[6] Harvey, *The Advocate's Devil*, 1985

[7] Court of Appeal, July 31, 1996

[8] More detail on the current position in English law can be found in Sommer, P, *Digital Footprints; Assessing Computer Evidence*, Criminal Law Review [1998] Crim LR (Special Edition) pp 61-78

[9] Frye v US 293 F.1013 (D.C. Cir. 1923)

[10] Daubert v. Merrell Dow Pharmaceuticals Inc 113 S.Ct. 2786 (1993) .

[11] Limpert P.Brad *Beyond the Rule in Mohan: A New Model for Assessing the Reliability of Scientific Evidence*, Toronto Law Review Vol 54 No 1. The *Daubert* principles extend into technical as well as scientific evidence: Kumho Tire Company, Ltd., et al, v. Patrick Carmichael, etc., et al. 97-1709 Supreme Court Of The United States

[12] The General Accounting Office Report GAO/AIMD-96-84 Defense Information Security and the Testimony of Jim Christy, Air Force Investigator, Senate Governmental Affairs Committee Permanent Investigations Sub-Committee, "Security in Cyberspace", June 5 1996, provide a description and give an idea of how the US authorities viewed the case. At http://www.bogus.net/kuji is a site created by Pryce's co-accused.

[13] There are a number of companies providing forensic tools directed at preserving data on PC hard-disks; they include Computer Forensics Ltd, Vogon/Authentec and Sydex. But adequate sector-by-sector copying is also possible using generally-available utilities such as PowerQuest's DriveImage (with smart-sectors switched off)

[14] For a general review of some of the problems, see Sommer, P *Evidence from Cyberspace: Downloads, Logs and Captures* Journal of Financial Crime 5JFC2 pp 138-152

[15] Schneier, B and Kelsey. J *Cryptographic Support for Secure Logs on Untrusted Machines*, The Seventh USENIX Security Symposium Proceedings, USENIX Press, January 1998, 53-62