Australian Computer Emergency Response Team

Collecting Electronic Evidence After a System Compromise

Matthew Braid, AusCERT, 2001

Collecting forensic evidence for the purposes of investigation and/or prosecution is difficult at the best of times, but when that evidence is electronic an investigator faces extra complexities. Generally, electronic evidence has none of the permanence that conventional evidence has, and is more difficult to present in a way that can be readily understood. The purpose of this paper is to highlight these difficulties and to suggest strategies to overcome them. Note that no legal advice is given here – different regions have different legislation. This paper will not address everything you need to know for your particular circumstances – it is a guide only. Always seek further information, including legal advice, for your specific circumstances.

Obstacles

Electronic crime is difficult to investigate and prosecute – often investigators have to build their case purely on any records left after the transactions have been completed. Add to this the fact that electronic records are extremely (and sometimes transparently) malleable and that electronic transactions currently have fewer limitations than their paper-based counterparts and you get a collection nightmare.

Computer transactions are fast – they can be conducted from anywhere, through anywhere, to anywhere; they can be encrypted or anonymous and generally have no intrinsic identifying features such as handwriting and signatures to identify those responsible. Any 'paper trail' of computer records they may leave can be easily modified or destroyed or may exist only temporarily. Worse still, auditing programs may automatically destroy the records left when they are finished with them.

Because of this, even if the details of the transactions can be retained or restored it is very difficult to tie the transaction to a person. Identifying information such as passwords, PIN numbers, or any other electronic identifier will not prove who did it – it merely shows that the attacker knew or was able to defeat those identifiers. Currently there is nothing that can be considered a true electronic signature for the purpose of criminal law in the same way that DNA or fingerprints do for other criminal investigations.

Even though technology is constantly evolving, investigating electronic crimes will always be more difficult due to the ability to alter data easily and because transactions may occur anonymously or deceptively. The best you can do is follow the rules of evidence collection as assiduously as possible.

Why Collect Electronic Evidence?

Given these obstacles, why bother collecting the evidence in the first place? There are two main reasons – future prevention and responsibility.

Future Prevention

Collecting electronic evidence involves investigating how the attack occurred. Without knowing what happened an organisation remains vulnerable to this type of attack and has little hope of stopping further attacks (including from the original attacker). It would be analogous to being defrauded for a large sum of money and not bothering to determine how the fraud was perpetrated. Even though the cost of collection can be high, the cost of repeatedly recovering from compromises is much higher, both in monetary and corporate image terms.

Responsibility

There are two responsible parties after an attack – the attacker and the victim. The attacker is responsible for the damage done and the only way to bring them to justice, to seek recompense and to deter further attacks is to convict them with adequate evidence to prove their actions.

Victims also have an ethical, if not legal, responsibility to the community. Sites that have been compromised and used to launch attacks against third parties may find that they – not the attacker – are sued for liability for the attack. The grounds for such a lawsuit might be that by failing to comply with the accepted minimum standards in network security they acted negligently. Public companies have a particular responsibility to their shareholders to ensure that business continuity and data confidentiality and integrity are not compromised. Victims may also have a legal obligation to perform an analysis of evidence collected, for instance if the attack on their system was part of a larger attack. For ethical reasons, some victims may see merit in sharing information gathered after a compromise with others to prevent further attacks.

Collection Options

Once a compromise has been detected you have two options – pull the system off the network and begin collecting evidence or leave it online and attempt to monitor the intruder. Both have their advantages and disadvantages. Monitoring may accidentally alert the intruder and cause them to wipe their tracks, destroying evidence as they go. If you disconnect the system from the network you may later find that you have insufficient evidence or, worse that the attacker left a 'dead man switch' that destroys any evidence once the system detects that it is offline. How you respond should be based on the situation. The "Collection and Archiving" section below contains information on what to do in each case.

Types of Evidence

Before you start collecting evidence it is important to know the different types of evidence categories. Without taking these into consideration you may find that the evidence you've spent several weeks and quite a bit of money collecting is useless.

Real Evidence

Real evidence is any evidence that speaks for itself without relying on anything else. In electronic terms, this can be a log produced by an audit function, provided that the log can be shown to be free from contamination.

Testimonial Evidence

Testimonial evidence is any evidence supplied by a witness. This type of evidence is subject to the perceived reliability of the witness, but as long as a witness is considered reliable, testimonial evidence can be useful and almost as powerful as real evidence. Written statements by a witness can be considered testimonial as long as the author is willing to state that they wrote it.

Hearsay

Hearsay is any evidence presented by a person who was not a direct witness. Written statements by someone without direct knowledge of the incident are hearsay. Hearsay is generally inadmissible in court and should be avoided.

The Five Rules of Evidence

In order for evidence to be considered useful, it must have the following properties:

1. Admissible

   This is the most basic rule – the evidence must be able to be used in court or elsewhere. Failure to comply with this rule is equivalent to not collecting the evidence in the first place, except the cost is higher.

2. Authentic

   If you can't tie the evidence positively to the incident, you can't use it to prove anything. You must be able to show that the evidence relates to the incident in a relevant way.

3. Complete

   It's not enough to collect evidence that just shows one perspective of the incident. Not only should you collect evidence that can help prove the attacker's actions but for completeness it is also necessary to consider and evaluate all evidence available to the investigators and retain that which may contradict or otherwise diminish the reliability of other potentially incriminating evidence held about the suspect. Similarly, it is vital to collect evidence that eliminates alternative suspects. For instance, if you can show the attacker was logged in at the time of the incident, you also need to show who else was logged in and demonstrate why you think they didn't do it. This is called Exculpatory Evidence and is an important part of proving a case.

4. Reliable

   Your evidence collection and analysis procedures must not cast doubt on the evidence's authenticity and veracity.

5. Believable

   The evidence you present should be clear, easy to understand and believable by a jury. There's no point presenting a binary dump of process memory if the jury has no idea what it all means. Similarly, if you present them with a formatted version that can be readily understood by a jury, you must be able to show the relationship to the original binary, otherwise there's no way for the jury to know whether you've faked it.

Using these five rules, we can derive some basic dos and don'ts.

1. Minimise Handling/Corruption of Original Data

   Once you've created a master copy of the original data, don't touch it or the original itself – always handle secondary copies. Any changes made to the originals will affect the outcomes of any analysis later done to copies. You should make sure you don't run any programs that modify the access times of all files (such as tar and xcopy), remove any external avenues for change and in general analyse the evidence *after* it's been collected.

2. Account for Any Changes and Keep Detailed Logs of Your Actions

   Sometimes evidence alteration is unavoidable. In these cases it is absolutely essential that the nature, extent and reasons for the changes be documented. Any changes at all should be accounted for – not just data alteration, but physical alteration of the originals (for instance the removal of hardware components) as well.

3. Comply with the Five Rules of Evidence

   The five rules are there for a reason. If you don't follow them you are probably wasting your time and money. Following these rules is essential to guarantee successful evidence collection.

4. Do Not Exceed Your Knowledge

   If you don't fully understand what you are doing, then it will be more difficult to account for any changes you make and you may not be able to describe what exactly you did. If you find yourself out of your depth and if time is available learn more before continuing otherwise find someone who knows the territory. Never soldier on regardless – you will just damage your case.

5. Follow Your Local Security Policy and Obtain Written Permission

During the course of your investigation you may be required to access and copy sensitive data or obtain statements from system users in which case there will be staff management issues to consider. Before commencing your investigation, it is important to ensure you have obtained written and signed permission to proceed and have clear instructions as to the scope of your investigation. Without clear authority to proceed, your actions may be, or be perceived to be, in breach of your company's security policy and you may find yourself personally accountable as a result. If in doubt, talk to those that know, including obtaining the necessary legal advice.

It is also recommended that your organisation develop appropriate policies and procedures for collecting electronic evidence so that they are in place *prior* to an incident occurring. This will significantly stream line the process and save valuable time before evidence is lost.

6. Capture as Accurate an Image of the System as Possible

This is related to point 1 – differences between the original system and the master copy count as a change to the data. You must be able to account for the differences.

7. Be Prepared to Testify

If you're not willing to testify about the evidence you have collected, you might as well stop before you started. Without the collector of the evidence being there to validate the documents created during the evidence collection process it becomes hearsay and inadmissible. Remember that you may need to testify at a later time.

8. Ensure Your Actions are Repeatable

No one is going to believe you if they can't replicate your actions and reach the same results. This also means that your plan of action shouldn't be based on trial-and-error.

9. Work Fast

The faster you work, the less likely the data is going to change. Volatile evidence (see below) may vanish entirely if you don't collect it in time. This is not to say you should rush – you must still collect accurate data and keep a record of your actions as you go. If multiple systems are involved, work on them in parallel (a team of investigators would be handy here), but each single system should still be worked on methodically. Automation of certain tasks makes collection proceed even faster.

10. Proceed From Volatile to Persistent Evidence

Some electronic evidence is more volatile than others. Because of this, you should always try to collect the most volatile evidence first.

11. Don’t Shutdown Before Collecting Evidence

    You should never shutdown a system before you collect the evidence. Not only will you lose volatile evidence but the attacker may have trojaned the startup and shutdown scripts, Plug-and-Play devices may alter the system configuration and temporary file systems may be wiped. Rebooting is even worse because it may result in further loss of evidence and should be avoided at all costs. As a general rule, until the compromised disk is finished with and restored it should never be used as a boot disk.

12. Don’t Run Any Programs on the Affected System

    Since the attacker may have left trojaned programs and libraries on the system, you may inadvertently trigger something that could change or destroy the evidence you’re looking for. Any programs you use should be on read-only media (such as a CD-ROM or a write-protected floppy disk), and should be statically linked.

Volatile Evidence

Not all the evidence on a system will last for extended periods of time. Some evidence resides in storage (i.e. volatile memory) only while there is a consistent power supply; other evidence stored is continuously changing. When collecting evidence, always try to proceed from most volatile to least volatile and from most critical to least critical machines/systems. For example, don’t waste time extracting information from an unimportant machine’s main memory when an important machine’s secondary memory hasn’t been examined.

To determine what evidence to collect first, draw up an Order of Volatility – a list of evidence sources ordered by relative volatility. An example Order of Volatility would be:

1.	Registers and Cache	6.	Main Memory
2.	Routing Tables	7.	Temporary File Systems
3.	Arp Cache	8.	Secondary Memory
4.	Process Table	9.	Router Configuration
5.	Kernel Statistics and Modules	10.	Network Topology

Once you have collected the raw data from volatile sources you may be able to shutdown the system.

General Procedure

When collecting and analysing evidence there is a four-step procedure you should follow. Note that this is a very generic outline – it may be necessary to customise the procedures to suit your situation.

Identification of Evidence

You must be able to distinguish between evidence and junk data. For this purpose you should know what the data is, where it is and how it is stored. Once this is done you will be able to determine the best way to retrieve and store any evidence found.

Preservation of Evidence

The evidence found must be preserved as close as possible to its original state. Any changes made during this phase must be documented and justified.

Analysis of Evidence

The stored evidence must then be analysed to extract the relevant information and to recreate the chain of events. Always be sure that the people who are analysing the evidence are fully qualified to do so.

Presentation of Evidence

Communicating the meaning of your evidence is vitally important – otherwise you can't do anything with it. It should be technically correct, credible and easily understood by persons with a non-technical background. A good presenter can help in this respect.

Collection and Archiving

Once you've developed a plan of attack and identified the evidence that needs to be collected, it's time to start capturing the data. Storage of that data is also important as it can affect how the data is perceived.

Logs and Logging

You should be running some kind of system logging function. It is important to keep these logs secure and to back them up periodically. Since logs are usually automatically timestamped a simple copy should suffice, although you should digitally sign and encrypt logs that are important to protect them from contamination. Remember that if the logs are kept locally on the compromised machine they are susceptible to alteration or deletion by an attacker. Having a remote syslog server and storing logs in a 'sticky' directory can reduce this risk, although it is still possible for an attacker to add decoy or junk entries into the logs.

Regular auditing and accounting of your system is useful not only for detecting intruders but also as a form of evidence. Messages and logs from programs such as Tripwire can be used to show what an attacker did. Of course, you need a clean snapshot for these to work, so there's no use trying it after the compromise.

Monitoring

Monitoring network traffic can be useful for many reasons – you can gather statistics, watch for irregular activity (and possibly stop an intrusion before it happens) and trace where an attacker enters and what they do.

Monitoring logs as they are created may show important information that might subsequently be deleted by the attacker. This doesn't mean that reviewing the logs later is not worthwhile – it may be what's *missing* from the logs that is suspicious.

Information gathered while monitoring network traffic can be compiled into statistics to define normal behaviour for your system. These statistics can be used as an early warning of an attacker's presence and actions.

You can also monitor the actions of your users. This can, once again, act as an early warning system – unusual activity (such as unsuccessful attempts to su to root) or the sudden appearance of unknown users warrants closer inspection.

No matter the type of monitoring done, you should be very careful – there are plenty of laws you could inadvertently break. In general you should limit your monitoring to traffic or user information and leave the content unmonitored unless the situation necessitates it. You should also display a disclaimer stating what monitoring is done when users log on. The content of this should be worked out in conjunction with your lawyer.

Methods of Collection

There are two basic forms of collection – 'freezing the scene' and 'honeypotting'. The two aren't mutually exclusive – you can collect frozen information after or during any honeypotting.

Freezing the scene involves taking a snapshot of the system in its compromised state. The necessary authorities should be notified (for instance the police and your incident response and legal teams) but you shouldn't go out and tell the world just yet.
You should then start to collect whatever data is important onto removable non-volatile media in a standard format and make sure that the programs and utilities used to collect the data is also collected onto the same media as the data. All data collected should have a cryptographic message digest created and those digests should be compared to the original for verification.

Honeypotting is the process of creating a replica system and luring the attacker into it for further monitoring. A related method – sandboxing – involves limiting what the attacker can do while still on the compromised system so they can be monitored without much further damage. The placement of misleading information and the attacker's response to it is a good method for determining the attacker's motives. You must make sure that any data on the system that refers to the attacker's detection and actions should be either removed or encrypted; otherwise they can cover their tracks by destroying it. Honeypotting and sandboxing are extremely resource intensive, so may be infeasible to perform. There are also some legal issues to consider, most importantly entrapment. As before – obtain legal advice.

Artefacts

Whenever a system is compromised, there is almost always something left behind by the attacker – be it code fragments, trojaned programs, running processes or sniffer log files. These are known as artefacts. They are one of the important things you should be collecting, but you must be careful. You should never attempt to analyse an artefact on the compromised system. They could do anything and you want to make sure their effects are controlled.

Artefacts may be difficult to find.Trojaned programs may be identical in all obvious ways to the originals (file size, MAC times etc). Use of cryptographic checksums may be necessary to determine whether files have been modified, so you may need to know the original file's checksum. If you are performing regular File Integrity Assessments, this shouldn't be a problem.

Analysis of artefacts can be useful in finding other systems the attacker (or their tools) has broken into.

Collection Steps

We now have enough information to build a step-by-step guide for the collection of the evidence. Once again this is only a guide – you should customise it to your specific situation.

1. Find the Evidence

   Determine where the evidence you are looking for is stored. Use a checklist – not only does it help you to collect it, but it can be used to double-check that everything you are looking for is there.

2. Find the Relevant Data

   Once you've found the evidence, you must identify what is relevant to the case. In general you should err on the side of over-collection, but you must remember that you have to work fast.

3. Create an Order of Volatility

   Now that you know exactly what to gather, work out the best order to gather it. Following the Order of Volatility for your system ensures that you minimise loss of uncorrupted evidence.

4. Remove External Avenues of Change

   It is essential that you avoid alterations to the original data. Preventing tampering with the evidence helps you to create as exact an image as possible, although you have to be careful, if you disconnect the system from the network, the attacker may have left a dead man switch. In the end you should try and do as much as possible.

5. Collect the Evidence

   You can now start to collect the evidence using the appropriate tools for the job. As you go, re-evaluate the evidence you've already collected. You may find that you missed something important. Now is the time to make sure you get it.

6. Document Everything

   Your collection procedures may be questioned later, so it is important that you document everything that you do. Timestamps, digital signatures and signed statements are all important – don't leave anything out!

Controlling Contamination – The Chain of Custody

Once the data has been collected it must be protected from contamination. Originals should never be used in forensic examination – verified duplicates should be used. This not only ensures that the original data remains clean, but also enables examiners to try more 'dangerous', potentially data-corrupting tests. Of course, any tests done should be done on a clean, isolated host machine – you don't want to make the problem worse by letting the attacker's programs get access to a network.

A good way of ensuring data remains uncorrupted is to keep a Chain of Custody. This is a detailed list of what was done with the original copies once they were collected. Remember that this will be questioned later on, so document everything. Record who found the data, when and where it was transported (and how), who had access to it and what they did with it. You may find that your documentation ends up greater than the data you collected, but it is necessary to prove your case.

Analysis

Once the data has been successfully collected it must be analysed to extract the evidence you wish to present and to rebuild what actually happened. As for other procedures, make sure you fully document everything you do – your work will be questioned and you must be able to show that your results are consistently obtainable from the procedures you performed.

Time

To reconstruct the events that led to your system being corrupted you must be able to create a timeline. This can be particularly difficult when it comes to computers – clock drift, delayed reporting and differing time zones can create confusion in abundance. One thing to remember is to never change the clock on an affected system. Record any clock drift and the time zone in use as you will need this later, but changing the clock just adds an extra level of complexity that is best avoided.

Log files usually use timestamps to indicate when an entry was added and these must be synchronised to make sense. You should also use timestamps – you're not just reconstructing events, you are contributing to the chain of events that must be accounted for as well. It's best to use the GMT (UTC) time zone when creating your timestamps – the incident may involve time zones other than your own, so using a common reference point will make things much easier.

Forensic Analysis of Back-Ups

When analysing backups, it is best to have a dedicated host for the job. This examination host should be secure, clean (a fresh, hardened install of the operating system is a good idea), and isolated from any network – you don't want it tampered with while you work and you don't want to accidentally contaminate others.

Once this system is available, you can commence analysis of the backups. Making mistakes at this point shouldn't be a problem – simply restore the backups again if required.

Remember the mantra – document everything you do. Ensure that what you do is not only repeatable, but that you always get the same results.

Reconstructing the Attack

Now that you have collected the data, you can attempt to reconstruct the chain of events leading to and following the attacker's break-in. You must correlate all the evidence gathered (which is why accurate timestamps are critical) – so it's probably best to use some graphical tools, diagrams and spreadsheets. Include all of the evidence you've found when reconstructing the attack – no matter how small it is. You may miss something if you leave a piece of evidence out.

As you can see, collecting electronic evidence is no trivial matter. There are many complexities to consider and you must always be able to justify your actions. It is far from impossible though – the right tools and knowledge of how everything works is all you need to gather the evidence required.

References

1. Collie, Byron S. “Intrusion Investigation and Post Intrusion Computer Forensic Analysis”. 2000.
   URL: http://www.usyd.edu.au/su/is/comms/security/intrusion_investigation.html

2. Collie, Byron S. “Collecting and Preserving Evidence after a System Compromise”. 2000.
   URL:
   http://mangrove.nswrno.net.au/dist/public/auugsec2000/Collecting%20and%20Preserving%20Evidence%20after%20a%20System%20Compromise.ppt

3. Romig, Steve. “Forensic Computer Investigations”. 2000
   URL: http://www.net.ohio-state.edu/security/talks/2001-10_forensic-computer-investigations/

4. McKemmish, R. (Australian Institute of Criminology) “What is Forensic Computing?” June 1999.
   URL: http://www.aic.gov.au/publications/tandi/ti118.pdf

5. Brezenski, Dominique andKillalea, Tom (Internet Engineering Task Force). “Guidelines for Evidence Collection and Archiving” July 2000.
   URL: http://www.globecom.net/ietf///draft/draft-ietf-grip-prot-evidence-01.html

6. Action Group into the Law Enforcement Implications of Electronic Commerce. “Issues Paper: Evidence and the Internet” September 2000.
   URL: http://www.austrac.gov.au/publications/agec/

7. Wright, T. “An Introduction to the Field Guide for Investigating Computer Crime (Part 1)” 17 April 2000.
   URL: http://www.securityfocus.com/infocus/1244

8. Wright, T. “The Field Guide for Investigating Computer Crime: Overview of a Methodology for the Application of Computer Forensics (Part 2)” 26 May 2000.
   URL: http://www.securityfocus.com/infocus/1245

9. Wright, T. “The Field Guide for Investigating Computer Crime: Search and Seizure Basics (Part 3)” 28 July 2000.
   URL: http://www.securityfocus.com/infocus/1246

10. Wright, T. “The Field Guide for Investigating ComputerCrime : Search and Seizure Planning (Part 4)” 1 September 2000.
    URL: http://www.securityfocus.com/infocus/1247

11. Wright, T. “The Field Guide for Investigating Computer Crime: Search and Seizure Approach, Documentation, and Location (Part 5)” 10 November 2000.
    URL: http://www.securityfocus.com/infocus/1248

12. Wright, T. “The Field Guide for Investigating Computer Crime, Part 6: Search and Seizure - Evidence Retrieval and Processing” 8 January 2000. URL: http://www.securityfocus.com/infocus/1249

13. Wright, T. “The Field Guide for Investigating Computer Crime, Part 7: Information Discovery - Basics and Planning” 26 February 2001. URL: http://www.securityfocus.com/infocus/1250

14. Wright, T. “The Field Guide for Investigating Computer Crime, Part 8: Information Discovery - Searching and Processing” 21 March 2001. URL: http://www.securityfocus.com/infocus/1251