



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Collecting Electronic Evidence After a System Compromise

Matthew Braid
SANS Security Essentials
GSEC Practical Assignment
Version 1.2b

Evidence is difficult to collect at the best of times, but when that evidence is electronic an investigator faces some extra complexities. Electronic evidence has none of the permanence that conventional evidence has, and is even more difficult to form into a coherent argument. The purpose of this paper is to point out these difficulties and what must be done to overcome them. Not everything is covered here – it should be used as a guide only, and you should seek further information for your specific circumstances. Note that **no legal advice is given here** – different regions have different legislation. If in doubt, always ask your lawyer – that's what they're there for.

Why Collect Evidence?

Electronic evidence can be very expensive to collect – the processes are strict and exhaustive, the systems affected may be unavailable for regular use for a long period of time, and analysis of the data collected must be performed. So why bother collecting the evidence in the first place? There are two simple reasons – future prevention and responsibility.

Future Prevention

Without knowing what happened, you have no hope of ever being able to stop someone else (or even the original attacker) from doing it again. It would be analogous to not fixing the lock on your door after someone broke in. Even though the cost of collection can be high, the cost of repeatedly recovering from compromise 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 (and to seek recompense) is with adequate evidence to prove their actions.

The victim on the other hand has a responsibility to the community. Information gathered after a compromise can be examined and used by others to prevent further attacks. They 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.

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 pros and cons. In the case of monitoring, you may accidentally alert the intruder while monitoring and cause them to wipe their tracks

any way necessary, destroying evidence as they go. You also leave yourself open to possible liability issues if the attacker launches further attacks at other systems from your own. If you disconnect the system from the network you may 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's offline. What you choose to do should be based on the situation. The "Collection and Archiving" section below contains information on what to do for either case.

Obstacles

Electronic crime is difficult to investigate and prosecute – investigators have to build their case purely on any records left after the transactions have 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), can be encrypted or anonymous, and 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 be only temporary. 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 restored through analysis it is very difficult to tie the transaction to a person. 'Identifying' information such as passwords or PIN numbers (or any other electronic identifier) does not prove who did it – it merely shows that whoever did it knew or could get past those identifiers.

Even though technology is constantly evolving, investigating electronic crimes will always be more difficult due to the ease of alteration of the data and the fact that transactions may be done anonymously. The best you can do is follow the rules of evidence collection and be as assiduous as possible.

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 the witness can be considered reliable, testimonial evidence can be almost as powerful as real evidence. Word processor documents written by a witness may 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. Word processor documents written by someone without direct knowledge of the incident is hearsay. Hearsay is generally inadmissible in court, and should be avoided.

The Rules of Evidence

There are five rules of collecting electronic evidence. These relate to five properties that evidence must have to be useful.

1. Admissible

This is the most basic rule - the evidence must be able to be used - in court or otherwise. 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 with 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 prove the attacker's actions but also evidence that could prove their *innocence*. 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 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 clearly understandable 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, human understandable version, 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 guaranteeing successful evidence collection.

4. Do Not Exceed Your Knowledge

If you don't understand what you are doing you can't account for any changes you make and you can't describe what exactly you did. If you ever find yourself out of your depth, either go and learn more before continuing (if time is available) or find someone who knows the territory. Never soldier on regardless - you're just damaging your case.

5. Follow Your Local Security Policy

If you fail to comply with your company's security policy you may find yourself with some difficulties. Not only may you end up in trouble (and possibly fired if you've done something *really* against policy), but you may not be able to use the evidence you've gathered. If in doubt, talk to those that know.

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 to 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 be collecting accurate data. 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 (see below) is more volatile than others are. Because of this, you should always try to collect the most volatile evidence first.

11. Don't Shutdown Before Collecting Evidence

You should never, ever shutdown a system before you collect the evidence. Not only do you lose any 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 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 is going to last very long. Some evidence is residing in storage that requires a consistent power supply; other evidence may be stored in information that is continuously changing. When collecting evidence, you should always try to proceed from most volatile to least. Of course you should still take the individual circumstances into account – you shouldn't waste time extracting information from an unimportant/unaffected machine's main memory when an important/affected machine's secondary memory hasn't been examined.

To determine what evidence to collect first, you should 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 general four -step procedure you should follow. Note that this is a very general outline – you should customise the details 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 work out the best way to retrieve and store any evidence you find.

Preservation of Evidence

The evidence you find 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. Analysis requires in -depth knowledge of what you are looking for and how to get it. Always be sure that the person or 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. The manner of presentation is important, and it must be understandable by a layman to be effective. It should remain technically correct, and credible. 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 the actual process of 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 any 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 damage 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 out for irregular activity (and possibly stop an intrusion before it happens) and trace where an attacker is coming from and what they are doing.

Monitoring logs as they are created can often show you important information you might have missed had you seen them separately. This doesn't mean you should ignore logs later – 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 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 should be considered definite cause for 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 related to the attacker's detection and actions should be either removed or encrypted; otherwise they can cover their tracks by destroying it. Honey potting and Sandboxing is extremely resource intensive, so it may be infeasible to perform. There are also some legal issues to contend with, most importantly entrapment. As before – consult your lawyers.

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, 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 figure out what of it 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 – don't spend hours collecting information that is obviously useless.

3. Create an Order of Volatility

Now that you know exactly what to gather, work out the best order to gather it. The Order of Volatility for your system is a very good guide as following 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, and prevention is always better than a cure. Preventing anyone from tampering with the evidence helps you to create as exact an image as possible, although you have to be careful – the attacker may have been smart and left a dead-man switch. In the end you should try and do as much as possible to prevent changes.

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 – who found the data, when and where it was transported (and how), who had access to it and what they did with it – everything. 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 everything you must 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, ever 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 in 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 be using timestamps – you're not just reconstructing events, you yourself are making a chain of events that must be accounted for as well. It's best to use the GMT time zone when creating your timestamps – the incident may involve other time zones than your own, so using a common reference point can 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 send something nasty down the line.

Once this system is available, you can commence analysis of the backups. Making mistakes at this point shouldn't be a problem – you can 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 you have 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 you must 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: ftp://ftp.net.ohio-state.edu/users/romig/other_papers/intrusion%20investigation.doc
2. Collie, Byron S. "Collecting and Preserving Evidence after a System Compromise". 2000.
URL: <http://mangrove.nswrno.net.au/dist/public/augsec2000/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/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 and Killalea, Tom (Internet Engineering Task Force). "Guidelines for Evidence Collection and Archiving" July 2000.
URL: <http://www.globecom.net/ietf/draft/draft-ietf-gnip-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/age_c/
7. Wright, T. "An Introduction to the Field Guide for Investigating Computer Crime (Part 1)" 17 April 2000.
URL: <http://www.securityfocus.com/focus/ih/articles/crimeguide1.html>
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/focus/ih/articles/crimeguide2.html>
9. Wright, T. "The Field Guide for Investigating Computer Crime: Search and Seizure Basics (Part 3)" 28 July 2000.
URL: <http://www.securityfocus.com/focus/ih/articles/crimeguide3.html>
10. Wright, T. "The Field Guide for Investigating Computer Crime : Search and Seizure Planning (Part 4)" 1 September 2000.
URL: <http://www.securityfocus.com/focus/ih/articles/crimeguide4.html>
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/focus/ih/articles/crimeguide5.html>

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/focus/ih/articles/crimeguide6.html>
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/focus/ih/articles/crimeguide7.html>
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/focus/ih/articles/crimeguide8.html>

© SANS Institute 2000 - 2002, Author retains full rights.