

Show and tell - digital forensics and giving evidence

Jonathan Haddock
October 2018



Please note that while there are a number of products linked to in this presentation, I'm not on commission for them nor do I have any investment in them. It is up to the reader to determine if a product meets their needs.



About me

- I've worked in IT for over 15 years
- Now look after the network & security for three local authorities
- Have run penetration testing courses with the YPISG
- Trained in forensics, penetration testing and incident response



<https://keybase.io/joncojonathan>



[@joncojonathan](https://twitter.com/joncojonathan)



<https://uk.linkedin.com/in/jonathanhaddock>

The YPISG is a BCS specialist group - Young Professionals' Information Security Group.

Every role helps with security, whether that's on the service desk or as a developer. I've worked across a number of roles, starting in service desk / desktop support, and each has benefited from experience gained in the others.

If you want to get into IT security, I recommend you specialise in an area you find interesting but make sure you keep your hand in other areas too so you have a broad skill set and knowledge. This will allow you to have informed discussions with different teams, in a way they (and you) can understand.

I also blog about cyber security, development, IT, and occasionally local history at <https://blog.jonsdocs.org.uk>.



What is digital forensics?

- Process of looking at computer systems and files to determine what activity took place
 - Can also be about recovery
- Increasingly important for incident response
- Not like it's shown in the movies!
 - You can't decrypt part of a drive like that
- Isn't "wet" forensics - we're in the digital realm
- Needs translation and explanation

Digital forensics ranges from the simple recovery of photos from a wiped SD card to a full investigation into the actions undertaken on a computer system.

There's an increased need for forensics awareness for incident response: determining how / where a cyber security incident (e.g. a malware infection or breach) occurred will allow you or others to contain the issue and prevent it happening again.

In the movies there's often a race against time to crack into the enemy's encrypted hard disk - note you can't usually break *part* of the key. Similarly, we can't enhance photos from poor pixelated imagery to HD level quality!

As this isn't "wet forensics" there's no DNA here, but we do have a number of advantages as I'll cover later.



Why translation?

- Logs can be really hard to read (or just plain horrible!)

```
188.65.57.4 - - [09/Oct/2018:20:32:40 +0000] "GET
/RSS/blogfeed.xml?blogID=1 HTTP/1.1" 301 185
"http://example.org/RSS/blogfeed.xml?blogID=1"
"SimplePie/1.2 (Feed Parser; http://simplepie.org; Allow
like Gecko) Build/20090627192103"
```

This just means someone accessed a file on a site, and got redirected (resource permanently moved)...

Part of the forensic analyst's job is to translate what they've found into something that can be understood. Logs in particular can be horrible to review and a judge / magistrate / HR manager is unlikely to understand the raw evidence you find.

This is an excerpt from a Nginx web server log. It's fairly unintelligible unless you know what you're looking at (Source IP, Date / time, HTTP method, Path, HTTP status code, URL, User agent, to name a few of the fields). Clearly we can't just present the logs to someone, we need to explain them.



What digital forensics isn't

- Not an instant indicator of blame
 - Also not your fault
- (Often) Unable to conclusively tie actions to a human being
 - Unless there's a camera pointing at the user and the screen, you can't know for certain who did what
- Not as easy as it used to be thanks to encryption

Because this isn't "wet forensics" it's virtually impossible to concretely say "this particular human being performed this particular action" - you'd need a camera looking at the keyboard, mouse, monitor and user in order to do that. You can, and I'd suggest should, state in your report that your evidence identifies a user account or device, not a human being. Digital forensic evidence can act as part of the overall evidence bundle, but shouldn't be used on its own to convict someone.

When performing an investigation you might find yourself thinking "I'm going to cost this person their job, my report is quite damning". This is a dangerous mindset that's easy to fall in to. What you need to remember is that *you* aren't going to cause someone to lose their job / go to jail, *their* actions have lead to this point (if they're guilty). It's important to be objective at all times and record your findings. Explain them but don't attach your own interpretation to them - that isn't your job as the analyst.

Don't forget we're not creating evidence, we're only reporting what's there.

Encryption has made investigation significantly harder, but it is good for privacy - a moral dilemma.

Encryption vs privacy

Tim Cook says Apple's refusal to unlock iPhone for FBI is a 'civil liberties' issue

- 'This is about much more than a single phone or a single investigation'
- FBI escalates war with Apple: 'Marketing bigger concern than terror'

Apple ordered to unlock San Bernardino gunman's phone




Dave Lee
North America technology reporter

17 February 2016

f + Share

San Bernardino shooting



Syed Rizwan Farook is seen in his California Department of Motor Vehicles photo

Apple has been ordered to help FBI investigators access data on a phone used by San Bernardino gunman Syed Rizwan Farook.

2016 saw the FBI ordering Apple to provide access to an iPhone used by a gunman. Apple refused, citing various reasons. The reason the FBI couldn't just look at the phone: it was encrypted.

More details:

- <https://www.theguardian.com/technology/2016/feb/22/tim-cook-apple-refusal-unlock-iphone-fbi-civil-liberties>
- <https://www.bbc.co.uk/news/technology-35593048>



Forensic process

- Obtain the assets
 - Files, disk images or physical media, phones
- Calculate the asset's hash
- Take an image of the asset
- Calculate hashes
- Make a copy
- Check its hash
- Perform analysis

If presented with files or pre-made images, it's wise to enquire about their provenance. How were they made, by whom, when? Is there a reason for not being provided the physical media / device?

We calculate a hash of the asset as this will describe the absolute original (*this was originally omitted from the presentation, added during question time, and has been added here for completeness*). Hashing is covered later in this presentation.

By taking an image of the asset we have a file we can perform our analysis on, protecting the original from changes. This is a great advantage over “wet forensics”. Consider the example of this room - with all of us having come to the event, any “wet forensic” evidence that was already in here has been damaged or destroyed by our presence. In digital forensics we can always refer to the original image, so our environment, so to speak, can always be pristine.

Every time the file moves we calculate the hash again, just to make sure the copy is valid.



Chain of custody

- First entry:
Assets received by you from provider
- Shows who had access, so possibility of tampering reduced
- If the asset moves, is copied or shown to someone else, note in the log
- Protects you

This document underpins all our work, and helps protect us / the case from allegations of tampering. As the log shows who could access the evidence it's possible to determine everyone that could have made changes or provided an interpretation.

Taking an image

- Not just a photograph!
- A block by block copy of a disk
 - Tools like `dd` or `ftkimg`
 - Not a copy of the files on the disk
- Means we can leave the original untouched
- Beware of Windows - it's not forensically safe
 - Use a write blocker



The difference between taking a copy of the *files* on physical media and making a *bit by bit* copy is the former only shows the current state of affairs. When a file is deleted, the file system (NTFS, EXT4, FAT32 etc.) simply marks it in the table as no longer being present. Space is also returned to the operating system so a new file could be placed where the old one was - but really the file is still there. In a *bit by bit* copy of the media we have a copy of the whole file system, not just the present files, so we can recover deleted files through a process known as *data carving*.

Microsoft Windows automatically mounts (makes available for use) disks and in doing so places a marker on the file system, thus changing the original media. Linux doesn't (always) do this, but in either case can be instructed to mount media read only (`mount -o ro`). If doing forensics on Windows it's important to use a write blocker (like the USB one, pictured) which intercepts and discards any instructions to write to the file system.



Taking hashes

- A hash is a mathematically calculated signature for a file
 - `da39a3ee5e6b4b0d3255bfef95601890afd80709`
- md5, sha1, sha2
- Note tool version in your report
 - `sha1sum (GNU coreutils) 8.29`
- The sum will change as soon as the file contents is edited

By taking or calculating a hash you determine a fingerprint for a file or device. As shown on the next slide, any change will produce a change to the hash.

After each file copy we calculate the hash of the copy to prove the files are the same.

Collisions are unlikely but can happen for MD5 and SHA1. A collision is when two different inputs provide the same hash. I used MD5 and SHA1 during my investigation as finding two different inputs that provide the same hash for both algorithms seems infeasible, especially when you consider the input would still have to be relevant to the case.



Hash example

- Simple text file, says “hello”

f572d396fae9206628714fb2ce00f72e94f2258f

- Text file edited to say “hell0”

40dfb682064a88af16ddb3faa2d3b5f0e6a82ee3

- A minor change has completely changed the hash

Example showing the SHA1 hash for a text file containing the text “hello”, and how it changes completely following the edit.



Make a copy

- Immediately after the original is received or the disk image created
 - Making disk images takes time, you don't want to do it again
- Check the hash
- Move the original to protected storage
- We'll work on a copy
- You may have to provide a copy

As mentioned previously, this is about protecting the original. After we've taken our image of the original asset it should be locked away, and noted in the chain of custody as such, with all work performed on a copy of the image.

Why on a copy? Images can take a long time to produce, and you don't want to repeat that exercise (especially given a lot of forensics is performed in environments where cost is very much a factor, and time is money).

You may be asked to provide a copy to expert witnesses or the opposition, so they can perform their own analysis.



Forensic tools

- Tools like *Encase* or *Forensic Tool Kit (FTK)* are expensive - around \$3,000 each
- *The Sleuth Kit* is open source
 - Has a graphical front end - *Autopsy*
- *PhotoRec* (free) can recover files
- But sometimes you just have to use your everyday applications

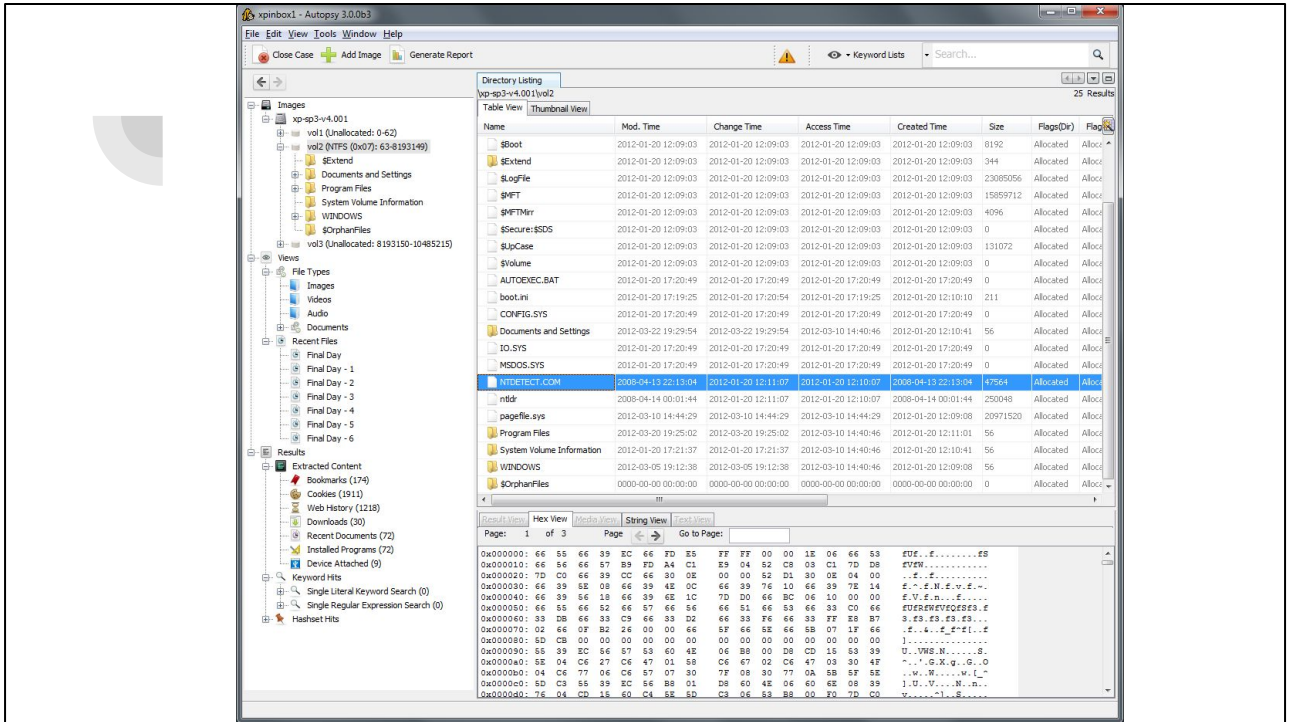
It's unlikely you'll have *FTK*

(<https://accessdata.com/products-services/forensic-toolkit-ftk>) or *Encase* (<https://www.guidancesoftware.com/>) to hand, as they're used by law enforcement or consultancies that undertake forensic work regularly. They are powerful tools, but for casual or irregular use perhaps a poor investment.

The Sleuth Kit (<http://www.sleuthkit.org/>) and *Autopsy* (<http://www.sleuthkit.org/autopsy/>) are free and becoming increasingly powerful. Initially *Autopsy* was a web based front end but as you'll see in the screenshot later on it's changed to be an application now. *Autopsy* supports a number of features, including the ability to look at browser histories.

Photorec by CG Security (<https://www.cgsecurity.org/wiki/PhotoRec>) is another free tool that started out as a data carver to recover photos, but now supports recovery of multiple file formats.

In my case I was dealing with PST file exports from a Microsoft Exchange server. Microsoft Outlook was my tool of choice as it's designed to work with these files, although I still followed forensic practice.



A screenshot of Autopsy, from its website.



Validate your tools

- As tools receive updates their output can change
- Have a known sample
- After an upgrade, confirm the tool still gives known results
 - If file `spoons.jpg` was found before, it still should be, and be identical
- New versions may add new features, more results

When a tool receives an update it's possible that it will behave differently. To determine this it's important to run tests against a known sample. Getting more information is not a problem, as the product may have simply added new features. No longer finding information would be a problem, and should be addressed through the vendor's support channels.



Analysing the evidence

- Work on a copy (check the hash)
- Pick the relevant tool
 - Cost
 - Task
- Stick to the specification
 - If looking for files or events in one date range, don't get distracted by others
- Try to avoid changing the copy

It's very easy to get distracted by your findings. In my case, I was reviewing a mailbox that contained thousands of emails, the vast majority of which were of no relevance to the investigation I was undertaking. I had no need, or right, to review those emails. To avoid distraction, stick to the brief: if you're looking for changes in a certain time frame, restrict yourself to it. You may find items relating to the same topic outside the time range, but they're outside of scope.

Avoid making changes wherever possible, even down to rotating pictures or renaming files in the image as changes can lead to confusion: did I do that, or was the file already like that? While we have the option of referring to a pristine copy, it's not desirable to have to keep looking at fresh copies as this takes time.

Each item we extract from the image (called an *artefact*), in my case emails, should have a hash calculated and this should be noted.



Keeping track

- I always keep a journal
 - Details of assignment
 - Findings (each artefact has a hash)
 - Screenshots
 - Thoughts
 - Chain of custody
- Journal should be shared on a need-to-know basis

Your journal is not what you present as your evidence, it's simply your working document. I record thoughts and to-do lists in mine, along with keeping a log of my chain of custody. Some forensic software allows you to record your notes in it, providing you a journal facility. If there were multiple researchers then the chain of custody would need to be a separate document.

How you structure your journal is up to you - whatever works best. When working on a forensic investigation my journal contains details of who authorised the investigation, a copy of the specification and many tables of data along with screenshots of assets. Conversely, when I'm performing a penetration test (ethical hacking) my journal looks different and includes spider diagrams. There are some common elements between the two, but I change based on my task.

Do whatever works best for you (if that's pen and paper that's fine, and I do that too, but remember to keep the physical copies secured).



Writing reports and statements

- Take information from your journal / notes
- Write for the audience - probably not too technical
 - Internal or external?
- Appendix:
 - List of hashes
 - Maybe chain of custody
- Also need-to-know

Your report or statement has to be *accurate*, that's the most important thing. Don't make stuff up, do your best to avoid errors, and don't be ambiguous. I read my statements through several times before I signed them (in my case, every single page needed my signature).

It may be appropriate to include your chain of custody, dependant on your audience.

The level of technical detail needs to be appropriate. Don't go into extensive detail on why data carving works - the reader isn't likely to need that information. Instead, adopt a style that allows you to remain brief and factual. You shouldn't be making any assumptions in your report - the facts will speak for themselves.

Refer to your journal for details to include in the report.

For me, the intended audience was barristers and eventually magistrates in court.

The image features a solid teal rectangular background. In the top-left corner, there is a cluster of five overlapping circles of varying shades of teal, arranged in a pattern that suggests depth. In the bottom-right corner, there is a series of four vertical bars, each composed of three overlapping circles of varying shades of teal, creating a sense of height and depth.

Court



Attending court

- You're not allowed in before giving your evidence
- Expect to sit around a lot
 - Courtroom might be triple booked
- Look at your questioner
- Answer to the Magistrates
 - Watch their pens - don't talk too fast
- Make sure your phone is off!

I was at a magistrates' court, meaning I was presenting evidence to three people that were acting like judges. They don't have a legal background, instead taking advice from a legal advisor.

My biggest worry was that I'd be discredited by the defence, like you see in the movies. My second biggest worry was that I had no idea what to expect in terms of questions. The barrister on your side (prosecution in my case) isn't allowed to tell you the questions in advance as that's known as "coaching a witness". Initially I wasn't allowed my statement for reference either, although the magistrates approved for me to receive this once I'd started giving evidence.

Strangely, I wasn't requested to provide any ID on arrival at court. I suspect this is because on first entering the witness box you're asked to swear an oath to tell the truth, and then state your full name. Lying under oath is committing perjury and has legal ramifications.

The magistrates may be taking notes - if so, watch their pens. Don't talk too fast as if you're too quick they won't be able to make notes (or will stop listening in order to do so).

Incidentally, you address the group of magistrates as "your worships".



Be clear, explain

- Magistrates are lay people
- Your job is to explain what the evidence means
- You might be allowed a copy of your statement
 - Can't be annotated
- If you don't understand the question, get it repeated or rephrased

In my case my technical evidence was not disputed, my attendance at court was to explain what it meant in a manner non-technical folk could understand. I could refer to my statement but importantly this wasn't annotated, so any dates I needed to confirm I had to find at the time.

To be fair to the defendant, I made sure to explain my investigation couldn't pinpoint a particular human being. In the UK we have the concept of being innocent until proven guilty and given my job as a witness is to assist justice, and present the truth, it was only fair to highlight that.



Don't make stuff up!

- You either found it in the evidence, or you didn't
- If you don't know, say you don't know
 - If you can't remember, say that too
- Answers should clarify or explain your statement
 - You're not creating new information, just explaining what's already presented

Remember - you're not on trial, but the court is interested to hear what you have to say. That said, if you make something up, or lie, you are committing an offence. As you're discussing your statement you don't necessarily need to worry about recalling details you've left out although you may be asked to explain part of the process.

You either found something that answers the question, or you didn't. You're not in court to create new information.



Cross examination

- The opposition barrister may want to ask questions
- Their aim is to help their client, don't feel attacked
 - Consider their questions as just other questions
- You may be asked for clarification to your earlier answers
- Alternatively, there may be entirely new questions
- The first barrister may ask additional questions

Cross examination worried me but wasn't a big deal. The questions were, literally, just other questions. Some questions I was asked were clearly intentionally ambiguous, so I ensured my answers were clear and precise.

After cross examination the magistrates asked me further questions. Treat them exactly the same: with respect.



Assumption trap

- “Is it fair to assume that...”
- Beware - is it *fair* to assume anything in criminal proceedings?
- Ultimately we’re not there to speculate
- Can be a way to use your expertise to cause you to lend support to a tenuous idea

I’d been warned about assumption questions by my employer’s legal team - be careful as this type of question can cause you to back yourself into a corner by making you think on your feet.

The important thing is that it’s *assumption*.

I was asked “is it fair to assume that because X happened on this date, that it’s the reason for Y”. I started looking at the dates, which matched, before realising this was an assumption trap. I pointed out that it wasn’t fair to assume anything in legal proceedings, as I had to be fair to the defendant and thus had to comment on my evidence.



Reporters

- There's likely to be reporters somewhere
 - Courtroom (public gallery)
 - Waiting area (public space)
 - Just outside
 - May be in disguise
- Don't be afraid to say "no comment"
- They can be pushy

I have no problem with the free press, so long as they're being honourable and accurate. During my time at court there were some reporters in the waiting area (they sat behind me), in the public gallery and outside. The one outside was in disguise (holding an unlit cigarette, I assumed he was going to ask me for a light).

Don't be afraid to say "no comment". As the magistrates pointed out to me before I left the courtroom, I wasn't to talk about the case to anyone outside (hence I've not gone into details here).

If representing your employer they may have their own policy on speaking to the media.

Any questions?

There were some good questions at the event, some summarised here.

When working on your investigate you mentioned you felt part of a team with colleagues in the ICO. Did you feel really isolated when giving evidence?

Ultimately, yes! While I could see the people I'd worked with (they were in court) I wasn't allowed to discuss the case with them until we'd all given evidence. For safety's sake, that meant not talking to them much at all so there was no suggestion of collusion. At the end of the day I was in court to give evidence based on what I'd discovered, so I could take comfort in my statement (the detail of which wasn't challenged).

You mentioned MD5 and SHA1, which are considered weak for some applications. Why did you use these?

MD5 and SHA1 are considered weak due to collisions - the ability to obtain the same hash for two different inputs. That said, finding an input that produces a collision that works for both MD5 and SHA1 would be difficult. I used these algorithms as they were what I had available on the system, and I recorded hashes for both for each asset, but arguably you could use hash algorithms from the SHA2 suite.

Given accounts can be compromised or used by someone else, how can you ever say a particular person did something?

The caveat stands that as digital forensic analysts we can only ever state an *account* (or device) which performed an action - we're restricted to the digital realm. One of the attendees was a lady from the law faculty who was able to add that the conclusion of any case is based on a bundle of evidence, and not based on a single item. To that end it's a case of showing "beyond all reasonable doubt".



Interested in security? Try a podcast

Smashing Security


Humorous run down of security news for the week. With Graham Cluley (@gcluley) and Carole Theriault (@caroletheriault).



www.smashingsecurity.com

Security Now

A semi-technical dive into recent security news and vulnerabilities. Hosted by Leo Laporte and Steve Gibson (@sggrc).




www.twit.tv/sn

Smashing Security is the less technical of the two, with episodes lasting around 45 minutes. Most episodes there's also a guest.


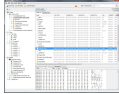
Security Now episodes tend to last around two hours.

**Thank you for
coming**





References and acknowledgements

- Screenshots of the San Bernardino case reporting:
 - <https://www.theguardian.com/technology/2016/feb/22/tim-cook-apple-refusal-unlock-iphone-fbi-civil-liberties>
 - <https://www.bbc.co.uk/news/technology-35593048>
- USB write blocker image: 
https://www.cru-inc.com/products/wiebetech/usb_writeblocker/
- Autopsy screenshot from 
<https://www.sleuthkit.org/autopsy/images/v3/overview.png>
- Podcast logos © their respective owners, from their respective sites