

Project 5: Forensics

This project is due on **Wednesday, April 29 at 6 p.m.** and counts for 8% of your course grade. Late submissions will be penalized by 10% plus an additional 10% every 5 hours until received. Late work will not be accepted after 20.5 hours past the deadline. If you have a conflict due to travel, interviews, etc., please plan accordingly and turn in your project early.

This is a group project; you will work in **teams of two** and submit one project per team. Please find a partner as soon as possible. If have trouble forming a team, post to Piazza's partner search forum.

Strict no-leaks policy. In this project you play the role of a computer forensic analyst working to solve a murder case. Since you don't want to be fired for jeopardizing an ongoing criminal investigation, you need to follow a strict policy on collaboration. *You are bound by the Student Code not to communicate with anyone regarding any aspect of the case or your investigation (other than within your group or with course staff).* The number of pieces of evidence you find, the techniques you try, how successful said techniques are, the general process you follow, etc. are all considered part of your solution and must not be discussed with members of other groups.

Start early. It may be impossible to complete this project before the deadline unless you begin several days beforehand. Please plan accordingly.

Solutions must be submitted electronically in any of the group member's svn directory, following the submission checklist below. Clearly state all member's netids of the team. Details on the filename and submission guideline is listed at the end of the document.

"In general, computer forensics is rather ad hoc. Traditional rules of evidence are broken all the time. But this seems like a pretty egregious example."

– Bruce Schneier

Introduction

In this project, you will play the role of a forensic analyst and investigate a murder mystery. On Thanksgiving 2009, a terrible crime occurred on University of Michigan campus. Hapless Victim, a leading figure in the university community, was shot while working in the EECS building. Detailed building information and map can be found on <http://campusinfo.umich.edu/campusmap/56>. Victim was last seen alive on November 25, 2009, shortly before midnight, and was discovered dead at approximately 6 a.m. the next morning. Officers recovered the projectile shown below (Exhibit A), which appears, inexplicably, to have been the cause of death.

The case went cold on December 1, when the leading suspect, Nefarious Criminal, fled the country and disappeared. Officers seized his computer, but the hard disk was encrypted and investigators were unable to crack the password. No further evidence could be found.



Exhibit A — Projectile recovered at the crime scene. Ballistics experts have identified it as a “Nerf blaster dart.”

There was finally a break in the case last week, when Nefarious Criminal was picked up by authorities on a beach in South America. Among his possessions was a Post-It note with the hard disk password on it. We have decrypted the drive and made it available for your analysis.

Your job is to conduct a forensic examination of the disk image and document any evidence related to the murder. If you find sufficient evidence, Nefarious Criminal will be extradited and face trial.

Objectives:

- Understand how computer use can leave persistent traces and why such evidence is often difficult to remove or conceal.
- Gain experience applying the security mindset to investigate computer misuse and intrusion.
- Learn how to retrieve information from a disk image without booting the operating system, and understand why this is necessary to preserve forensic integrity.

Getting Started

The tools and techniques you use for your investigation are up to you, but here are some suggestions to help you get started. NOTE: This MP requires several tools to be downloaded and installed. For live analysis, you **MUST** use your own machine to perform the analysis as Virtual Machine is not supported on EWS Linux system. Also, decompressed disk image and virtual hard disk image file you will be generating sums upto >7GB which will take majority of your EWS quota. However, for dead analysis, Autopsy tool is already installed on EWS Linux and can be performed on the system. Finally, we strongly recommend you to run password cracking softwares on your own machine.

General Knowledge A general working knowledge of Linux is undoubtedly helpful for this project. If you don't have this yet, you may need to spend time Googling and/or experimenting to get up to speed. The TA will also answer general Linux questions as a last resort. For an excellent reference book, try *UNIX and Linux System Administration Handbook* by Nemeth, Snyder, Hein, and Whaley. Ebook can be found at <https://books.google.com/books>. See http://en.wikipedia.org/wiki/Disk_partitioning for some additional background.

Live Analysis Live analysis is a forensic technique in which the investigator examines a running copy of the target system. We suggest using VirtualBox for this purpose.

1. Download the compressed raw disk image (1.2 GB):
`https://subversion.ews.illinois.edu/svn/sp15-cs461/_shared/mp5/forensics_release_w15.raw.gz`
2. Decompress the disk image.
3. Convert the raw disk image to a VirtualBox disk image:
`VBoxManage convertdd forensics_release.raw forensics.vdi -format VDI`
4. Use the VirtualBox GUI to create a new VM. Select Linux / Ubuntu (32-bit) as the machine type. Select "Use an existing virtual hard disk file" and select the VDI you just created.
5. Start the VM and explore the system.

Dead Analysis In dead analysis, the forensic investigator examines data artifacts from a target system without the system running. We suggest trying dead analysis with the Autopsy open-source forensics tool. The procedure below assumes you are working on Ubuntu Linux. (If you like, you can reuse the VM from the previous project.) Autopsy will also run on Windows and OS X. For OS X, Sleuth Kit should be installed prior to Autopsy installation. Sleuth Kit can be downloaded from <http://www.sleuthkit.org/sleuthkit/download.php>.

1. Install the Autopsy digital forensics suite:
`$ sudo apt-get install autopsy`

If you are using EWS machine, load sleuth kit:

```
$ module load sleuthkit
```

2. Launch Autopsy in the background and open the browser-based GUI:

```
$ sudo autopsy &
```

In a browser on the local machine, go to the URL <http://localhost:9999/autopsy>.

3. Create a new case and add the disk image:

- (a) Click New Case. Enter a case name and click New Case.

- (b) Go back to <http://localhost:9999/autopsy> and open the case you created.

- (c) Click Add Host. Enter a host name and click Add Host.

- (d) Click Add Image. Click Add Image File. Enter the path to the decompressed raw disk image. Make sure you select Type=Disk and Import Method=Symlink. Click Next.

- (e) Leave the Image File Details and File System Details as the defaults. (Note that the disk image contains 3 partitions, which Autopsy will allow you to examine separately.) Click Add. Click OK.

- (f) Select a partition to examine and click Analyze. The buttons at the top give you several analysis tools. Try File Analysis and Keyword Search to get started.

4. In addition to hints dropped elsewhere, here is an incomplete list of things to try:

- Examine the system logs.
- Check for deleted or encrypted files.
- Search the drive image for strings that may indicate relevance to your investigation.

Password Cracking Password crackers may be helpful in trying to brute-force decrypt password-protected files. John the Ripper (<http://www.openwall.com/john/>) is the canonical Unix password cracker. Hydra (<http://www.thc.org/thc-hydra/>) is a tool used to brute force remote login passwords, fcrackzip (<http://home.schmorp.de/marc/fcrackzip.html>) is a ZIP password cracker, and pdftcrack (<http://sourceforge.net/projects/pdftcrack/>) is a PDF password cracker. John, fcrackzip, and pdftcrack are conveniently available in the Debian package repositories and can be installed with `apt-get`.

When using a password cracker, it is wise to make sure that the password is not susceptible to a dictionary attack and does not use a restricted character set (e.g., lowercase letters, letters only, letters and numbers only) before spending time on a full brute-force crack. It is also a good idea to crack a very vulnerable password first to make sure you are using the tool correctly.

Tasks and Deliverables

The deliverables for this project are your answers to the questions below. Your answers should be *complete* but *concise*. None of the questions should require more than 1–2 paragraphs to answer.

For each prompt, explain the investigatory methods you used and the evidence that supports your conclusion. Place your responses in a plain text file called `report.txt`. If you recover files that are relevant to your responses, mention them by name and include them with your submission in a directory named `evidence/`.

1. Try booting the suspect's machine and using it normally. What *specific* behaviors of this machine make this a bad idea? Attach relevant evidence.
2. What operating system does the suspect use? Be careful and specific; e.g., say "Windows 2000" instead of just "Windows."
3. What is the username of the account typically used by the suspect?
4. Are there any indications that the suspect had an accomplice who was physically present on the night of the crime? Attach relevant evidence.
5. Were there any suspicious-looking encrypted files on the machine? If so, please attach them and their decrypted contents as evidence and briefly describe how you obtained the contents.
6. Are there any indications that the suspect owned or was researching weapons of the kind involved in the murder? Attach the specific evidence and give a brief explanation.
7. Did the suspect try to delete any files before his arrest? Please attach the name(s) of the file(s) and any indications of their contents that you can find. (Hint: We will be impressed if you manage to recover the *original* contents of a particular incriminating file, but we do not expect you to do so.)
8. Reconstruct the timeline of actions by the suspect that may be relevant to the investigation. (Make a list in this format: <date> <time>: <event description>.) Include any activities related to your other responses, if you can identify when they occurred. Include each time the suspect logged in to or booted the machine. When was the last activity before the suspect fled the country?
9. Is there anything else suspicious about the machine? Attach relevant evidence as appropriate.

As you investigate, be on the lookout for evidence of any other machines or network services that the suspect may have used. These may contain important evidence and raise further questions you'll need to investigate (*hint, hint!*). Before attempting to access any such machines or accounts, be sure to contact your supervisor for permission by posting *privately* on Piazza. Failure to ask permission is guaranteed to be a waste of your time and may violate the course ethics policy. Again, start early; headquarters has been known to take up to 24 hours to approve such requests.

Policies and Hints

Collaboration: Strictly prohibited outside your group. As stated above, you are bound by the Honor Code not to communicate with anyone regarding any aspect of the case or your investigation (other than within your group or with course staff). The number of pieces of evidence you find, the techniques you try, how successful said techniques are, the general process you follow, etc. are all considered part of your solution and must not be discussed with members of other groups. If anyone brings up the project, start yelling “LALALALA” and refer them to your supervisor for an official spokesperson.

If you get stuck... Requesting Hints Given the nature of this assignment and its strict collaboration policy, HQ recognizes the need for some hints. We have standard hints for each question. If your group gets stuck, post *privately* on Piazza with the names of your group members, the question for which you would like a hint, and the progress you have made thus far on that question. Use the SAME Piazza thread your team has created to ask follow-up questions.

To help you get started, hints on the first three questions are free (and so are requests to access remote machines). After that, each group may receive a maximum of 3 hints, and we will enforce a one-hour delay between hints for each group. Purely administrative questions or general questions about Linux do not count towards this limit. All the questions from same group should be posted in SINGLE Piazza thread for staffs to conveniently track the progress and the number of hints provided.

Submission Checklist

We have generated blank file and directory in the mp5 folder of your svn repository. Please edit the file and/or upload files in the requested folder.

<code>partners.txt</code>	A plain text with one netid on each line.
<code>report.txt</code>	A plain text file with your answers to the numbered prompts.
<code>evidence/</code>	A directory containing any recovered files referenced in your report.