Project 5: Forensics

This project is due on **Thursday, December 6** 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 you 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 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 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.

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 via Canvas, following the submission checklist below. Please coordinate carefully with your partner to make sure at least one of you submits on time.

Introduction

In this project, you will play the role of a forensic analyst and investigate the theft of company secrets from SuperDuperSketchyCorp (SDSC). SDSC became aware of the theft after **The Media** ran a story regarding one of their closley guarded secrets.

The case went cold on March 17th, when the leading suspect, Leslie Nielson, fled the country and disappeared. Officers seized their computer, but the hard disk was encrypted and investigators were unable to crack the password. No further evidence could be found.

Investigators just recently caught a break when they found the hard disk encryption password on a sticky note in Leslie's home office. They've decrypted the device 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 crime. If you find sufficient evidence, a case can be brought against Leslie.

Learning 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.

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. TA's 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. Also, 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 raw disk image (4.3 GB): https://eecs388.org/*/forensics_release.tar.gz
2. Decompress the disk image. Note: some browsers seem to automatically decompress after having downloaded the file. Check to see if the file size is about 10 GB; if so, it has been decompressed already. Otherwise, decompress the image.
3. Use the VirtualBox GUI to create a new VM. Select Linux / Ubuntu (64-bit) as the machine type. Select "Use an existing virtual hard disk file" and select the VDI you just downloaded.
4. 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 (with a great GUI!) and OS X. Autopsy is not mandatory but is available as a resource should you need it.

1. Convert the vdi disk image to a raw disk image:
   Note: This is a hard drive so extracting it to RAW makes the file about 30 GB in size.
   `VBoxManage internalcommands converttoraw forensics.vdi forensics_release.raw`

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

3. 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.

4. 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.

5. 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 pdfcrack (http://sourceforge.net/projects/pdfcrack/) is a PDF password cracker. John, fcrackzip, and pdfcrack 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 and corresponding evidence to the questions below. Your answers should be *complete* but *concise*. None of the questions should require more than 2–3 sentences to answer. Keep in mind you are writing a technical report. You are expected to present your findings in a clear and easy to read manner. Failure to do so may result in a formatting deduction of up to 5% of your grade for this project.

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 shared the company secrets with other individuals? Attach relevant evidence.

5. 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 to do something interesting. When was the last activity before the suspect fled the country?

6. Is there anything else on the computer that would imply the suspect had malicious intents in using this computer? Attach relevant evidence as appropriate.

7. Were there any suspicious-looking encrypted files on the machine? If so, please attach them and their decrypted contents (if possible) as evidence and briefly describe how you obtained the contents.

8. Did the suspect try to delete any files before their 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.)

9. List all other secrets, and other findings. (In the format of <finding>: <explanation of discovery>)

As you investigate, be on the lookout for evidence of any other machines or network services or websites 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, accounts or websites, be sure to contact your supervisor for permission by emailing eecs388-proj5@umich.edu. Failure to ask permission is guaranteed to be a waste of your time and may violate the course ethics policy as well as resulting in a deduction from the assignment grade. 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 or 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. If your group gets stuck, email eecs388-proj5@umich.edu 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. Each group is allotted 3 hints. Please note, that headquarters has been known to take up to 24 hours to get back to agents.

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.

Finishing the Project This project is intentionally open-ended. As such, it is your duty to follow all leads you can find, and to conduct a thorough investigation. Your evidence will be paramount in solving this case and putting the right person behind bars.

Headquarters is unable to answer questions regarding whether a report contains all of the possible findings, since it's still an ongoing investigation. It's your job as a forensic analyst to draw as complete of a picture of the crime as possible. You should submit when you believe you have conducted a thorough investigation.

Submission Checklist

Upload to Canvas a gzipped tar file named `project5.`*`uniqname1`*`.`*`uniqname2`*`.tar.gz` that contains only the files listed below. You can generate the tarball at the shell using this command:

`tar -zcf project5.`*`uniqname1`*`.`*`uniqname2`*`.tar.gz report.txt evidence/`

The tarball should contain only the files below:

`report.txt`	A plain text file with your answers to the numbered prompts.
`evidence/`	A directory containing any recovered files referenced in your report.