

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

Introduction

The Autopsy Forensic Browser is a graphical interface to utilities found in The Sleuth Kit (TSK). TSK is a collection of command line tools that allow you to investigate a Windows or Unix system by examining the hard disk contents. TSK and Autopsy will show you the files, data units, and metadata of NTFS, FAT, EXTxFs, and UFS file system images in a read-only environment. Autopsy allows you to search for specific types of evidence based on keywords, MAC times, hash values, and file types.

Autopsy is HTML-based and uses a client-server model. The Autopsy server runs on many UNIX systems and the client can be any platform with an HTML browser. This enables one to create a flexible environment with a central Autopsy server and several remote clients. For incident response scenarios, a CD with The Sleuth Kit and Autopsy can be created to allow the responder read-only remote access to a live suspect system from an HTML-browser on a trusted system.

Autopsy will not modify the original images and the integrity of the images can be verified in Autopsy using MD5 values. There are help pages for the main analysis modes and The Sleuth Kit Informer is a newsletter that adds additional documentation. This document provides an overview of how to use Autopsy and what it can do.

Installation

1. Install The Sleuth Kit. It will be the most convenient if symlinks are made from the specific version directory (i.e. sleuthkit-1.00) to a generic one (i.e. sleuthkit).
2. Untar the Autopsy file.
3. Run 'make'. It will try to locate the grep and strings utilities. If any are not found, it will prompt you for the location. It will also prompt for the directory where The Sleuth Kit was installed.
4. The install script will ask if you have the NIST National Software Reference Library (NSRL). If you do, you will need to enter the path of it. The NSRL is available from www.nsrl.nist.gov.
5. You will be prompted for the Evidence Locker location. This is the base directory where all cases will be stored. You must create this directory on your own.
6. Type 'make live' or run the 'make-live-cd' script to build the 'live-cd' directory. The 'live-cd' directory can be burned to a CD.

Case Management

Starting with Autopsy 1.70, you can have multiple cases. When Autopsy is started, there is an Evidence Locker directory (specified on the command line or at installation time). This directory is the base where all cases will be stored.

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

A CASE is any investigation and can have one or more hosts in it. A list of investigators is assigned to each case. Each case gets a subdirectory of the evidence locker and there is a configuration file for the case and the list of investigators.

A HOST is a subset of a CASE. A host contains one or more file system images that are analyzed. Each host gets a subdirectory in the case directory. Each host has its own configuration file that describes the files that it uses. Each host also has five directories in it:

- images: for all the disk and partition images - this should have strict permissions to prevent modification
- output: for all output files from tools. This includes unallocated disk space and data unit contents.
- logs: Audit logs and investigator notes are stored here
- reports: All ASCII and HTML reports can be stored here
- mnt: Can be used to mount the images in loopback mode

An IMAGE corresponds to a disk or partition image. Image files are imported into an Autopsy host. The image file must be a raw copy of a partition or disk. These can be created by the 'dd' tool. Issue 11 of The Sleuth Kit informer discussed how to make images using 'dd'.

When importing an image, you have the option of moving the image to the evidence locker, copying the image to the evidence locker, or making a symbolic link from its current location to the evidence locker. You also have the option to calculate or add the MD5 hash value of the image.

Main Functions

After you have setup your case and imported the file system images, you can begin the investigation. The Host Gallery view provides a list of the imported file system images and you can select one of them to analyze. After you have selected it, you will enter the analysis view. The top of the window will have a series of tabs that represent different analysis modes.

Each mode performs a different type of analysis. Choose the mode that will help you find the type of evidence you are looking for. If you are looking for a specific file, choose the File mode. If you have a specific keyword in mind, choose that mode. If you are looking for a specific file type, then choose that mode. You will now need to use your sleuthing skills to search for evidence. You may want to refer to some books dedicated to this topic if you have not done this before.

Modes

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

FILE BROWSING: Allows browsing the image as a file system. This gives a list of directories on the left, and files and file content on the right hand side. The output of each file can be seen as ASCII or can be run through strings.

Since this analyzes directory entries, deleted file names can still be seen and depending on the OS, the deleted file contents can also be easily recovered. If a file name has a check before it, it has been deleted. The directory contents listings can be resorted based on name, size, times etc. by selecting the proper column header.

KEYWORD SEARCHING: Search an image using grep for a given string. The result will be a list of data units that have this string. Selecting each unit brings the user into Data Unit mode to view the contents. Case insensitive searches and 'grep' regular expression searches can also be performed. To decrease the searching time, a file can be generated with just the ASCII strings of the image. Also, the unallocated data can be extracted and searched to make deleted data recovery more efficient.

The search.pl file contains predefined search values. Autopsy currently comes with a regular expression to identify date strings and IP addresses. Additional values can be added by the user. The format is given in the file.

TIMELINE ANALYSIS: A timeline of file activity can be created and viewed. The timeline allows one to identify file and directory locations to examine. The times associated with files can be easily modified, so the time line should be used as reference only.

IMAGE DETAILS: Details about the file system are displayed. Examples of this mode include the Volume name, last mount time, and the physical layout of the data structures. For FAT file systems, the FAT contents are given and UNIX-based systems show the group layouts.

FILE TYPE ANALYSIS: Data reduction is an important aspect of digital forensics. One way of doing data reduction is to exclude known files and identify unknown files or categories. The File Type Analysis mode will examine all of the files in an image and sort them based on their file type. For example, all JPEG and GIF files would be identified as 'images'. This mode can also identify files that have an extension that is different than its file type. This uses the 'sorter' tool from The Sleuth Kit. The hash databases are used in this mode to exclude files that are known to be good and identify 'known bad' files. Refer to issues 3, 4, and 5 of The Sleuth Kit Informer for more details.

METADATA BROWSING: Metadata is descriptive data about a file. This includes information such as times, owner id, and a list of data unit pointers. This mode allows one to view the contents of the file system structures that hold these values. In UNIX-based file system these are typically called inodes, for FAT they are directory entries, and for NTFS they are MFT entries. In this mode, one enters the address of the structure and the details are shown. The file(s) that are using the file will also be displayed (even if they have been deleted for some OSes).

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

Metadata browsing can also be entered from within File browsing. When the file's metadata address is selected, the browser switches to metadata mode and displays the associated details. The data units that the file has allocated can be viewed using the data unit browsing.

DATA UNIT BROWSING: All file systems need to store file data some where. Typically, the file system space is organized into large chunks of consecutive bytes. These chunks have different names depending on the file system type, so we will just refer to them as data units. For UNIX-based file system the chunks are fragments, FAT are sectors or clusters, and NTFS are sectors.

This mode allows one to examine any data unit they want. Just enter the address and it is displayed in a variety of formats. This is most useful when used with searching or metadata browsing. The contents of the data unit can be displayed in ASCII, hexdump, or by running the raw output through strings(1). The metadata structure that has allocated the unit will be displayed (if any) along with the file name (if any).

There are two types of data unit addresses in Autopsy, regular and unallocated. The regular address is the unit number in a regular image created from dd. The unallocated address is the unit number in an image created from the unallocated units in a regular image (by using dls). When unallocated addresses are entered, they are converted to the regular address and the corresponding regular unit is shown. This is useful when using Autopsy along with foremost <http://foremost.sourceforge.net> or Lazarus (TCT).

INVESTIGATOR NOTES: An investigator can add notes about any file, data unit, or metadata structure. The notes can be viewed through Autopsy at the Main Menu or by any text editor. The notes file is saved in the 'logs' directory. When viewing through Autopsy, the location that the note refers to can be easily viewed.

REPORT GENERATION: Each of the above browsing techniques allows a report to be generated. This report lists the date, md5 value, investigator, and other context information in a text format. This can be used for record keeping when deleted data units of data have been found.

THE CELL: In an ideal world, forensics should only be performed on an air-gapped network. In some cases, such as incident response of critical systems, this is not possible. For this reason and because of a history of HTML-browser security issues, files in Autopsy are not "interpreted" by your browser. For example, an HTML document by default will be shown as the raw HTML text. If an investigator wants to view the actual HTML output or an image, they can do so in a sterilized environment that parses out embedded scripts and off-site references. Refer to issue #1 of The Sleuth Kit Informer for more details.

Regular Usage

1. Ensure that the evidence locker directory has been created and start Autopsy: `./autopsy`

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

Copy and Paste the URL into an HTML browser on the local system. It will look something like: ***http://localhost:9999/autopsy***

2. Select the 'Create Case' button and enter a name and list of valid investigator names. Note that both the case and investigator names must be valid directory names.
3. Select the case from the Case Gallery and then select 'Add Host' in the Host Gallery menu. Enter the host name, and time information such as the timezone and clock skew (if known). The timeskew is how many seconds fast or slow the original system was and the output times will be adjusted using it. For example, if the host was 3 seconds slow, this field would get a '-3'.
4. Select the host from the Host Gallery and then select 'Add Image' in the Host Manager menu. Copy the images to the directory shown on the screen. It is a subdirectory of the Evidence Locker for the new host and case that have been created. After the images are in the directory, press 'Refresh'. The images must be partition images in a raw format (i.e. dd).
5. Select the file system type and mounting information. By default, the MD5 value will be calculated for the image and saved for future integrity checks. If you already know it, select 'Add Known Value' and paste it in.
6. Continue to add images and hosts to the case. When done, select one of the images and using the different browsing modes.

Common Configurations

The basic usage is for a single user with the client and server on the same system. Autopsy 1.70+ can now handle more than one case at a time. The syntax is as follows for the server to run on port 9999 and only allow access from localhost:

```
# ./autopsy
```

To specify a different port number, use this:

```
# ./autopsy -p 8888
```

To specify a different remote host, use this:

```
# ./autopsy 10.0.0.1
```

To specify both a port and remote address use:

```
# ./autopsy -p 8888 10.0.0.1
```

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

If more than one investigator is going to be using the same server, then just choose different ports:

```
# ./autopsy -p 9000 10.0.0.1
```

and

```
# ./autopsy -p 9050 10.0.50
```

You can also specify a new evidence locker location by providing the '-d' argument:

```
# ./autopsy -d /usr/local/forensics2
```

Security Considerations

The Autopsy server is a Perl program that only processes Autopsy urls. It offers easy access control restrictions by limiting access to the server to one host and uses a random numeric "cookie" to further authenticate a remote user. The random cookie is generated when the server starts and allows an investigator to use a multi-user machine. The recommended usage is to have the browser and autopsy running on the same single-user system, which is the default behavior.

If a non-localhost system is specified, a cookie is automatically generated. If localhost is used, then a cookie is not used by default. The default behavior can be changed using the command line arguments. SSH forwarding can be used if encryption is needed over a network.

File names must be very simple (letters, digits, -, _, and .). This allows fast and easy checking of file names passed in the URL and does not allow people to move out of the morgue directory. Symbolic links can be created between the simple names and more complex ones.

Troubleshooting

1. Autopsy is complaining that it can't find X: Verify the variable settings in conf.pl.
2. Autopsy takes a very long time to display large directories: This occurs because directory contents are displayed as an HTML table, and many browsers are not very efficient at displaying large tables. So, it is not Autopsy that is slow, it is the browser.
3. Autopsy hangs when opening directories: Same answer as previous question. Browsers don't like big tables.
4. Autopsy is getting slower and slower: If you start an intensive operation, such as searching or making a strings file, and you hit the back button you will not stop the search or operation. There is no current way to stop these types of processes besides issuing a 'kill' command from a shell.

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

5. Errors are generated by the 'strings' and 'grep' utilities: This occurs because you most likely do not have the GNU version and the flags are not working. Install the GNU grep and binutils and verify that Autopsy is pointing to them in conf.pl.
6. Internet Explorer gives protocol and host errors: If you are accessing the localhost, then use the 127.0.0.1 IP address instead of the localhost name.
7. A file system image doesn't show up on the menu: Make sure your version of Perl supports large files.

Live Analysis

Live analysis is, in my mind, an investigation that occurs using the software resources of the suspect system. An example scenario of this is when a suspect system is found running, a CD is placed into it, and commands are run. If the suspect system is powered down and booted from a bootable Linux CD (or similar), then the investigation is a dead analysis.

This is most commonly done when investigating a server or other computer that is suspected of being compromised, but verification is needed before it can be powered down. Using The Sleuth Kit and Autopsy will prevent the access times on individual files from being updated (although the raw device's A-time will be) and can bypass most rootkits that hide files and directories.

Live analysis is not ideal because you are relying on the suspect system, which can lie, cheat, and steal. In addition to the potential of getting false information from the operating system you will also overwrite memory and maybe swap space during the investigation. If you are interested in examining the memory of the system, you should probably acquire that before you begin a live analysis.

An issue with doing live analysis with Autopsy is that it requires Perl, which is a large program and will likely need to depend on libraries and other files on the suspect system.

You will want to have a trusted CD for a live analysis, and autopsy makes that fairly easy. Compile autopsy as you would for a normal dead analysis installation. Then execute 'make live' in Autopsy. This script will make a 'live-cd' sub-directory in the autopsy directory, which contains a copy of autopsy and copies of TSK executables, grep, strings, perl etc:

```
# make live
Making base directory (./live-cd/)
Copying executables
Copying autopsy files
Creating configuration file using existing settings
```

Try the 'make static' with TSK to see if you can make static executables for your platform.

AUTOPSY INSTALLATION & USE

Taken From the Autopsy Documentation (Brian Carrier)

The 'live-cd' directory has a 'bin' directory where additional executables can be copied to and then the whole directory can be burned to a CD.

After the CD has been created and there is a system suspected of being compromised, then it is time to take advantage of the new features. There are two scenarios for live analysis. The first scenario uses a network share from a trusted system that you can write to. In this case, autopsy is run as normal and you specify the evidence locker directory as the mounted disk. The evidence locker is specified with '-d':

```
# ./autopsy -d /mnt/ev_lock 10.1.32.123
```

The above would start autopsy, use '/mnt/ev_lock/' as the evidence locker and would allow connections from 10.1.32.123 (where the investigator would connect from using an HTML browser). Remember that we do not want to write to the suspect system, so we should only use a network share and not a local directory in this scenario.

The second scenario does not use an evidence locker and does not intentionally write any data to disk. This scenario does not need the network share and each of the devices (or partitions) that will be analyzed are specified on the command line using the '-i' flags. The '-i' flag requires three arguments: the device, the file system type, and the mounting point. For example, to examine the '/dev/hda5' and '/dev/hda8' partitions on a Linux system, the following could be used:

```
# ./autopsy -i /dev/hda5 linux-ext3 / -i /dev/hda8 linux-ext3 /usr/10.1.32.123
```

The file system type must be one of the types that are supported by TSK. The remote IP address must also be given, otherwise you will have to use a browser on the suspect system and that will write data to the disk.

When you use the '-i' flag, then autopsy will start in the 'Host Manager' view where you can select the image that you want to analyze. You will skip the case and host configuration. The default case name will be 'live', the default host name is 'local', and the default investigator name is 'unknown'.