Microsoft Filesystems

Digital Forensics

Filesystem Content

- Filesystems hold files, primarily regular files which hold data, and directories (aka folders) which hold lists of files
- Directories are used to organize the files on the disk in a virtual hierarchy so that users can easily store and retrieve the data they want to work with
- Data contained in files also usually has an organization or structure which enables applications to store and retrieve relevant data for the application's use
- Files also have metadata which is information about a file, regardless of its specific contents, such as file type, ownership, permissions, timestamps, and potentially many other pieces of information about a file

Filesystem Structure

- The design of a filesystem determines what kinds of files it can hold, what metadata is stored for those files, and the limits on these things
- A filesystem resides in a volume, so a complete image of a volume includes all the filesystem metadata as well as the files
- There are 2 main filesystem formats (or designs) used for Microsoft systems, FAT (in several variations) and NTFS (also in more than one version)

FAT

- FAT (File Allocation Table) format is the traditional format for Microsoft operating systems and has been in use since the days of DOS
- It has been re-implemented in several versions over time and remains in use today (FAT12, FAT16, FAT32, exFAT, etc.)
- Windows NT replaced FAT with NTFS which addressed several significant shortcomings in the FAT format, notably increasing limits and adding metadata
- Many Windows computers have both formats in use

NTFS

- NTFS was created to overcome limitations in the FAT filesystem formats
- Main changes included
  - Increased maximum filesystem size, maximum file size, maximum file count
  - Added file attributes, streams, compression
  - Improvements to security, recoverability, scalability
- NTFS has also undergone 5 re-implementations
- NTFS v3.1 (commonly called NTFS5 or NTFS5.0) is the fifth version

What is out there on Windows

- Windows NT was the first OS to ship with NTFS and also used FAT32
- Windows CE 6.0 added exFAT to the mix, which is kinda sorta like FAT, but with licensing restrictions and intended for different devices than the typical computer disk, most notably as the default format for SDXC cards larger than 32GB
- Optical media uses filesystems designed for optical media, such as ISO9660
- There are some other filesystem types that can be used with a Windows computer if software is installed to use them
- Most Windows computers have both NTFS and FAT in use, as well as using optical media filesystems as required

Analysis Techniques

- Early digital investigators had to become experts in the low level details of filesystem data structures in order to extract anything that wasn't in plain sight on the disk
- Modern forensics tools do the hard work for us in normal situations, we no longer have to dig in the bits to get past the typical user efforts to hide data
- We now focus on identifying files of interest, and the metadata that gives context to those files - the who, what, when of the user's interactions with those files
- We browse, we search, and we use automated scanning to find files or interactions of interest
- It is possible for serious bad actors to hide information in normally unused areas of the filesystem such as gaps and in slack spaces but the hidden information is not that hard to find with modern forensic tools and steganography is much more secure, so these hiding methods are not expected to be found much anymore

What is of interest?

- Some items of interest are easy to identify (e.g. application files containing multimedia, business documents, financial statements, etc.)
- Some are part of application data stores (e.g. Outlook pst or ost files, excel spreadsheets, databases, etc.)
- Some are no longer present, but have left tracks we can find (e.g. registry data, search indexer data, application caches, hibernation files, etc.)
- Finding these context or trace items is best done with a tool that knows what to look for, although it can be done by hand by someone who is very knowledgeable about the trace and context information used by specific applications
- Paid software suppliers excel in this area because they can afford the resources to buy, build, test, and maintain the required software, documentation, training, and support staff

What Is Interesting About Files?

- The most interesting thing about a file is usually whatever it contains, i.e. a document or multimedia
- Other things are also relevant but not necessarily obvious
  - Where was it in the filesystem
  - Who owned it
  - When was it made, last changed, last accessed (sometimes called mac times)
  - What permissions, size, kind of file was it
  - Other FS-based metadata such as acls and streams
  - Is there metadata contained in the file such as URLs and application data
- Metadata may not be trustworthy (e.g. file came from another fs type)

Image Analysis Tools

- Most professional forensics software companies have been making significant changes in how they provide their software, and how they work with files from competitors's software
- There are open source no charge options, including the Autopsy package which runs on multiple platforms
- Some Linux distros include The Sleuth Kit, Autopsy and other no-charge forensics tools
- There are configurations of popular Linux distros to turn them into Forensics stations e.g. Caine, SIFT
- Live boot operating systems are useful for mobile forensics work

CLI Tools

- There are many CLI tools for forensics, and they have been around for quite a few years and are the underpinnings of many professional forensics packages
- They require a greater knowledge of filesystems than GUIs, but can be used in imaginative ways that GUIs are unsuitable for
- They allow for working with very large data volumes more efficiently than GUIs which make assumptions about what you might want to do and often preprocess things in ways you may not have any use for
- The resources list includes several of them

GUI Tools

- This is the area that professional tools rule
- A lot of the effort in paid product development goes into GUIs to try to appeal to a wider range of customers
- There are open source GUIs, and you can roll your own for specific tasks using the command line tools behind the scenes
- FTK Imager is a popular GUI tool for doing captures on Windows systems using Windows with a physical write blocking device
- Autopsy is a popular free GUI analyzer

Web Applications for Forensics

- Autopsy can run as a web app under Linux, allowing a team to use web browsers to work on a case database
- Web apps allow for remote or mobile work in ways GUIs do not, although if you need to capture large amounts of data, uploading over the internet is not usually feasible
- The autopsy team is not keeping the web app up to date and over the last few years have focused exclusively on the Windows GUI version
- The GUI version is the one to use for analysis on Windows

Sequence of Actions for Image Investigations

- Identify volume containing evidence, note in report
- Capture image, note in report
- Identify unusual volume attributes such as large gaps, note in report
- Identify filesystem(s) present on volume, note in report
- Identify files of interest in filesystem, tag or extract for report
- Include registry data and deleted items info as appropriate, tag or extract for report
- Analyze free space and unallocated space for any useful items if appropriate, extract for report
- Save evidence hashes for analyzed items of interest in report

Role of Forensics Software

- Depending on the software tools used, the investigator may need to create the report by hand, fill in components of the report, or include a generated report with their own summary document
- Generated reports often are highly technical and require translation to layman's terms to be useful
- Forensics software often only generates lists of the artifacts found, not what the artifacts represent

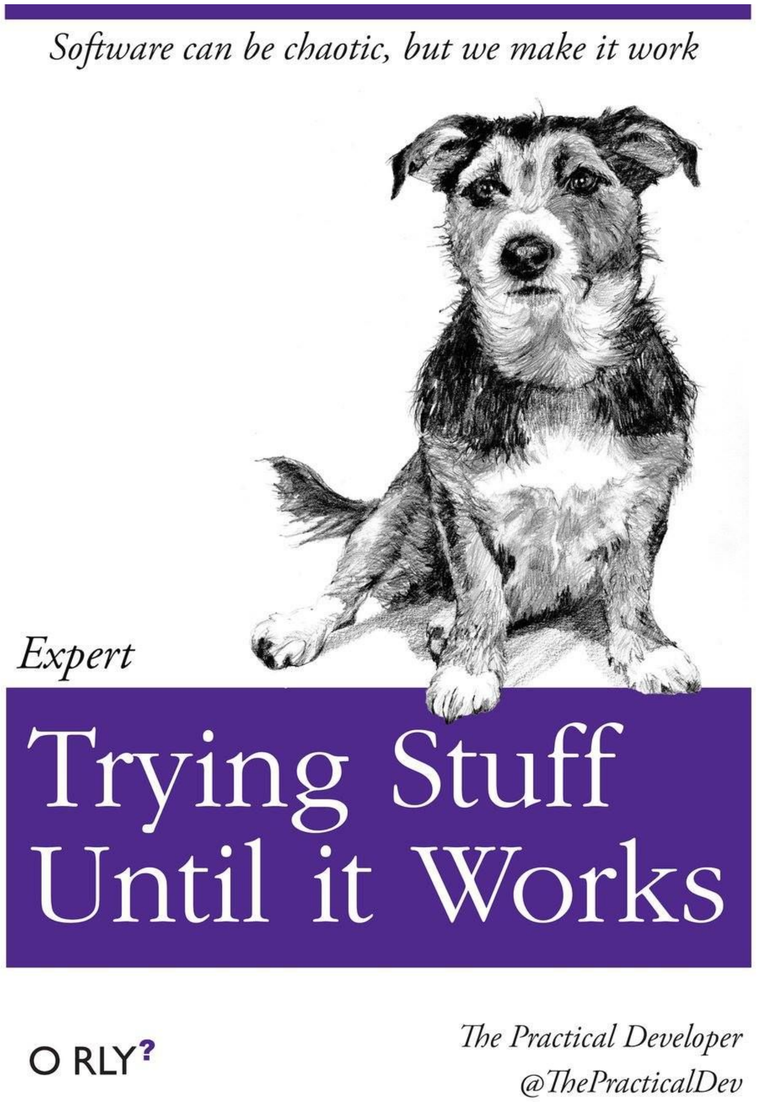

Microsoft Filesystem Analysis Lab

- Using Autopsy to find deleted files
- using ProDiscover to find evidence on a Windows drive image
- Using Autopsy to find relevant NTFS file stream data
- Using Access Data Registry Viewer to find relevant trace data in a registry file