Practical Computer Forensics using Open Source tools

Bruce Nikkel

June 12, 2008

Presentation Summary

- Overview of Digital Forensics
- Overview of Open Source Computer Forensic Tools
- Practical Examples
- Resources and Q&A

What is Digital Forensics?

- The collection, preservation, analysis, and presentation of digital evidence...
  - Admissible in a court of law
  - Usable for employee disciplinary hearings
  - Supporting data for internal incidents
  - Assisting/furthering other investigations

Field of Digital Forensics

- Computer forensics (hard disk, removable media acquisition and analysis)***
- Network forensics (network intrusions, abuse, etc.)
- Software forensics (examining malicious code, malware, etc.)
- Live system forensics (compromised hosts, system abuse, etc)

Digital Evidence is Data that...

- Helps reconstruct past events or activity (timelines)
- Shows possession/handling of digital data
- Show use/abuse of IT infrastructure & services
- Shows evidence of policy violation or illegal activity

Difficulties of Digital Evidence

- Easy to destroy
  - starting a PC updates hundreds of timestamps and modifies many files
  - attaching a hard disk or USB stick will modify file system timestamps
  - volatile memory is lost when a machine is powered off
- Hard to get
  - network traffic only exists on the wire for milliseconds
  - intrusions and attacks may be cleverly devised
  - anti-forensic activity may prevent collection

Overview of Open Source Computer Forensic Tools

- disk acquisition/imaging, and forensic image formats
- disk and file system analysis
- unallocated blocks, deleted files, slack-space recovery
- data carving
- "known good" hash databases

dcfldd

- Developed by U.S. Dept. of Defense Forensics Lab
- based on traditional Unix dd, but rewritten with forensics in mind
- cryptographic hashing for evidence preservation, error handling, logging, splitting, verification
- used for "forensically sound" disk acquisition

sleuthkit forensic suite

- Developed by Brian Carrier, based on original Coroner's Toolkit (TCT) by Farmer and Venema
- A set of analysis tools for getting info about:
  - disk layouts, partition tables (DOS, BSD, Sun, GPT)
  - filesystems, files, directories
  - timestamps and filesystem timelines
  - deleted files, unallocated areas, slack-space

foremost

- Devloped by Jesse Kornblum and Kris Kendall (U.S. Air Force Office of Special Investigations), based on scapel
- "data carving" forensic tool, attempts to extract files from unstructured data
- uses analysis of headers, footers, and known file formats
- useful for corrupt disks, swap, memory dumps, network traffic, or any "blob" of unknown data

Autopsy

- A side project of Sleuthkit, developed by Brian Carrier
- Web-based front-end for:
  - basic case management
  - analysis using sleuthkit tools

PyFlag

- FLAG (Forensic and Log Analysis GUI)
- Web GUI interface, and set of command line tools
- Analysis of:
  - disk devices and disk images
  - captured network traffic (pcap)
  - logs

NSRL Databases

- National Software Reference Library, maintained by NIST
- a database of hashes identifying files from known software packages
- Used to filter out "known good" files

Afflib and tools

- Forensic image format and aff tools, developed by Simson Garfinkel
- Intended to be an open, peer-reviewed, vendor independent standard
- Allows the direct working with compressed files (ie. allows seeking)
- Sleuthkit is compatible with AFF
- Hold other meta data about the image and case

Examples... disk acquisition

- taking an md5 hash during acquisition

```
dcfldd hash=md5 if=/dev/hda of=image.dd
```

- taking an sha1 hash of every 1Gb of the disk

```
dcfldd hash=sha1 hashwindow=1G if=/dev/hda of=image.dd
```

- verify a disk against an image:

```
dcfldd vf=/home/bruce/image.dd if=/dev/sdg
```

disk layout, partition table info

- mmls displays disk layout and partition scheme

```
# mmls /dev/sda
DOS Partition Table
Offset Sector: 0
Units are in 512-byte sectors

Slot    Start        End          Length       Description
00:  -----   0000000000   0000000000   0000000001   Primary Table (#0)
01:  -----   0000000001   0000000062   0000000062   Unallocated
02:  00:00   0000000063   0034298774   0034298712   Linux (0x83)
03:  00:01   0034298775   0035873144   0001574370   DOS Extended (0x05)
04:  -----   0034298775   0034298775   0000000001   Extended Table (#1)
05:  -----   0034298776   0034298837   0000000062   Unallocated
06:  01:00   0034298838   0035873144   0001574307   Linux Swap / Solaris x86 (0x82)
07:  -----   0035873145   0035888129   0000014985   Unallocated
```

File system info

- fstat displays much info about the filesystem

```
# fsstat /dev/sda1

FILE SYSTEM INFORMATION
--------------------------------------------
File System Type: Ext3
Volume Name:
Volume ID: 3d6c8a6fef240a9dc04def540921d90c

Last Written at: Mon Jun  9 22:00:08 2008
Last Checked at: Mon May 26 01:22:32 2008

Last Mounted at: Mon Jun  9 22:00:08 2008
Unmounted properly
Last mounted on:

Source OS: Linux
Dynamic Structure
Compat Features: Journal, Ext Attributes, Resize Inode, Dir Index
InCompat Features: Filetype, Needs Recovery,
Read Only Compat Features: Sparse Super, Has Large Files,

Journal ID: 00
Journal Inode: 8
```

file system info (cont.)

```
METADATA INFORMATION
----------------------------------------
Inode Range: 1 - 1073152
Root Directory: 2
Free Inodes: 889080
Orphan Inodes: 855140, 852843, 852841, 460954, 329085, 856376, 856660,
856659, 8
56658, 856657, 856654, 856653, 856652, 881066, 499732,

CONTENT INFORMATION
----------------------------------------
Block Range: 0 - 4287338
Block Size: 4096
Free Blocks: 3256977

BLOCK GROUP INFORMATION
----------------------------------------
Number of Block Groups: 131
Inodes per group: 8192
Blocks per group: 32768

Group: 0:
  Inode Range: 1 - 8192
  Block Range: 0 - 32767
  Layout:
    Super Block: 0 - 0
    Group Descriptor Table: 1 - 2
    Data bitmap: 1025 - 1025
    Inode bitmap: 1026 - 1026
    Inode Table: 1027 - 1282
    Data Blocks: 1283 - 32767
  Free Inodes: 8181 (99%)
  Free Blocks: 0 (0%)
  Total Directories: 2
```

File and directory analysis

- Listing all files: (recursive and full path):

```
fls -r -p partition.dd
```

- Listing just deleted files, or just regular files:

```
fls -r -p -d partition.dd
fls -r -p -u partition.dd
```

- Listing just directories or just files:

```
fls -D
fls -F
```

- Get more/long info (file_type inode file_name mod_time acc_time cre_time size uid gid) with -l

File and directory analysis

- Sample output (default and long):

```
d/d 817861:     var/lib/apt/mirrors
d/d 817878:     var/lib/apt/mirrors/partial
d/d 817862:     var/lib/apt/periodic
r/r 817865:     var/lib/apt/periodic/update-stamp
r/r * 817899(realloc):  var/lib/apt/extended_states.tmp

d/d 4579428:    tmp/dir1      2008.06.12 08:27:24 (CEST)      2008.06.12 08:27:22 (CEST)
      2008.06.12 08:27:24 (CEST)      4096    1000    1000
d/d * 4579438:  tmp/dir2      2008.06.12 08:27:29 (CEST)      2008.06.12 08:27:12 (CEST)
      2008.06.12 08:27:29 (CEST)      0       1000    1000
r/r * 2488828:  tmp/file1.txt 2008.06.12 08:23:30 (CEST)      2008.06.12 08:23:11 (CEST)
      2008.06.12 08:23:30 (CEST)      0       1000    1000
r/r 2488829:    tmp/file2.txt 2008.06.12 08:23:25 (CEST)      2008.06.12 08:23:25 (CEST)
      2008.06.12 08:23:25 (CEST)      406     1000    1000
r/r * 2488671(realloc): tmp/test.txt      2006.04.27 12:57:05 (CEST)
2007.01.09 00:48:59 (CET)       2007.01.08 16:03:55 (CET)       463     0       0
```

Recovery of deleted files

- locate file from the fls output, get inode number
- extracting file from the inode:

```
icat -o 4193280 image.dd 700740 > filename
```

- This could be a normal file or a deleted file
- "-s" includes the slackspace of the file

Recovery of (un)allocated and slack space

- extracting allocated space:

```
dls -a partition.dd > alloc.dls
```

- extracting unallocated space:

```
dls -A partition.dd > unalloc.dls
```

- extracting slack space:

```
dls -s partition.dd > slack.dls
```

- readable output: "-a" ascii or "-h" hex, add "-w" to view them as html

```
dcat -h image.dd 5436
```

- Use dcalc to map extracted data back to the image

Filesystem timelines

- mactime creates a log style timeline of each timestamp on each file
- use fls with the -m flag to prepare data for mactime (a prefix directory must also be specified)
- can be piped from fls directly into mactime:

```
fls -r -m / partition.dd | mactime -b -
```

- Just show a certain date range:

```
mactime -b timeline.data 01/21/2004-01/27/2004
```

Filesystem timelines (Cont.)

- Example mactime output:

```
Wed Jun 11 2008 23:32:53   4096  m.c  d/drwxr-xr-x   0      0      409601   /media
                          12288  m.c  d/drwxr-xr-x   0      0      458753   /etc
Wed Jun 11 2008 23:54:20   4096  m.c  d/drwxrwxrwt   0      0      499713   /tmp
Thu Jun 12 2008 00:44:30   4096  .a.  d/drwxrwxrwt   0      0      499713   /tmp
                           4096  m.c  d/drwxr-xr-x   0      0      335873   /root
Thu Jun 12 2008 00:44:47    142  mac  -/-rw-r--r--   0      0      27894    /document.txt
Thu Jun 12 2008 00:44:56      0  .a.  -/-rw-r--r--   0      0      27895    /test (deleted)
Thu Jun 12 2008 00:45:02     33  .a.  l/lrwxrwxrwx   0      0      27448    /initrd.img.old
```

Carving unstructured data

- List possible files for extraction in an image:

```
foremost -wv -i image.dd
```

- Extract all known files in an image, and save to a sorted directory:

```
foremost -t all -i image.dd
```

- Extract all jpegs:

```
foremost -t jpeg -i image.dd
```

NSRL Database samples

- NSRL Database format, one line for each file:

"001A5E31B73C8FA39EFC67179C7D5FA5210F32D8","49A2465EDC058C975C0546E7DA07CEE","E93AF649","CNN01B9X.GPD",83533,8762,"Vista",""

"000C89BD70552E6C782A4754536778B027764E14","0D3DD34D8302ADE18EC8152A32A4D934","7A810F52","gnome-print-devel-0.25-9.i386.rpm",244527,2317,"Linux",""

"001A6684A98A452F8501CD6F2D4A287A8FD5B709","F6F49036001D752F6F378247D911018D","7C46DD00","AppleTalk.h",78184,2490,"MacOSX10.2",""

"0067CB46B52B6ABEB5FC6362D7B4791021537C46","DA23D20200F82E94AECDAA4D37F169D6","096FE2DC","NETWATCH.EX_",16869,524,"WIN311",""

Resources

- Web resources
  - www.e-evidence.info, a directory of digital forensics documentation and papers
  - www.forensicswiki.org, a Wikipedia style forensics website
  - www.forensicfocus.com, an online forensics community
- Books
  - File System Forensic Analysis, Brian Carrier
  - Forensic Discovery, Dan Farmer, Wietse Venema
- Peer reviewed practitioner/research journals
  - Elsevier's Digital Investigation Journal, The International Journal of Digital Forensics & Incident Response
  - International Journal of Digital Evidence (IJDE)

Questions? Comments?

- Questions or comments?
- Contact me at nikkel@digitalforensics.ch
- Slides available at www.digitalforensics.ch