



Security of Cloud Computing

Fabrizio Baiardi
f.baiardi@unipi.it



Syllabus

-
- Cloud Computing Introduction
 - Definitions
 - Economic Reasons
 - Service Model
 - Deployment Model
 - Supporting Technologies
 - Virtualization Technology
 - Scalable Computing = Elasticity
 - Security
 - New Threat Model
 - New Attacks
 - Countermeasures
- ← Incident Handling and Detection

Technical Definition: Digital Forensics

“Tools and techniques to recover, preserve, and examine digital evidence on or transmitted by digital devices.”





Definition for the Masses

“Deleted” information, on almost any kind of digital storage media, is almost never completely “gone”...

Digital Forensics is the set of tools and techniques to recover this information in a forensically valid way (i.e., acceptable by a court of law)



Motivation

Deleted files aren't securely deleted

- Recover deleted file + when it was deleted!

Renaming files to avoid detection is pointless

Formatting disks doesn't delete much data

Web-based email can be (partially) recovered directly from a computer

Files transferred over a network can be reassembled and used as evidence



Motivation (2)

Uninstalling applications is much more difficult than it might appear...

“Volatile” data hangs around for a long time (even across reboots)

Remnants from previously executed applications

Using encryption properly is difficult, because data isn't useful unless decrypted

Anti-forensics (privacy-enhancing) software is mostly broken

“Big” magnets (generally) don't work

Media mutilation (except in the extreme) doesn't work

- **Basic enabler: Data is very hard to kill**



Traditional Digital Forensics Investigation

What's possible?

- Recovery of deleted data
- Discovery of when files were modified, created, deleted, organized
- Can determine which storage devices were attached to a specific computer
- Which applications were installed, even if they were uninstalled by the user
- Which web sites a user visited...



Traditional (2)

What's not...

- If digital media is completely (physically) destroyed, recovery is impossible
- If digital media is securely overwritten, recovery is very, very complicated, or impossible

Privacy Through Media Mutilation



degausser

or



or

or



forensically-secure
file deletion
software
(but make sure it works!)



Who Needs It?

Law enforcement

- Prosecution of crimes which involve computers or other digital devices
- **Defend the innocent**
- Prosecute the guilty
- Must follow strict guidelines during entire forensics process to ensure evidence will be admissible in court

Military

- Prosecution of internal computer-related crimes
- Own guidelines, many normal legal issues do not apply



Who (2)

General

- Employee misconduct in corporate cases
- What happened to this computer?
- For accidental deletion or malicious deletion of data by a user (or a program), what can be recovered?
- Need for strict guidelines and documentation during recovery process may or may not be necessary

Privacy advocates

- What can be done to ensure privacy?
- Premise: Individuals have a right to privacy. How can individuals ensure that their digital data is private?
- Very difficult, unless strong encryption is used, then storage of keys becomes the difficult issue



Digital Forensics: Goals (1)

Identification of potential digital evidence

- Where might the evidence be? Which devices did the suspect use?

Preservation of evidence

- On the crime scene...
- First, stabilize evidence...prevent loss and contamination
- Careful documentation of everything—what's hooked up, how it's hooked up...
- If possible, make identical, bit-level copies of evidence for examination



Digital Forensics: Goals (2)

Careful extraction and examination of evidence

- Directory and file analysis

Presentation of results of investigation (if appropriate)

- “The FAT was fubared, but using a hex editor I changed the first byte of directory entry 13 from 0xEF to 0x08 to restore ‘HITLIST.DOC’...”
- “The suspect attempted to hide the Microsoft Word document ‘HITLIST.DOC’ but I was able to recover it by correcting some filesystem bookkeeping information, without tampering with the file contents.”

Legal: Investigatory needs meet privacy



Digital Forensics: Constraints

Order of volatility

- Some data is more volatile
- RAM > swap > disk > CDs/DVDs
- Idea: capture more volatile evidence first

Chain of custody

- Maintenance of possession records for all
- Must be able to trace evidence back to original source
- “Prove” that source wasn’t modified = compute hash

Legal Issues



Admissible in court?

- Generally yes, but there is limited precedent.
- Shooting a moving target. But if it is consistent and no evidence gets created, it *should* be OK.

Legal to gather?

- Yes (with an appropriate court order) and yes for certain other circumstances
- Often it is the *only* way to gather information, as the court order may specify that machines cannot be taken down
- Network sniffing is considered a wire tap. Be careful!
 - requires a court order
 - often hard to get
 - same for incoming text messages on cell phones(!)



Investigatory Process: Needs

Acceptance

- Steps and methods are accepted as valid

Reliability

- Methods can proven to support findings
- e.g., method for recovering an image from swap space can be shown to be accurate

Repeatability

- Process can be reproduced by independent agents



Investigatory (2)

Integrity

- Evidence is not altered (if at all possible) and can prove that was not altered (or measure the degree to which it was altered)

Cause and effect

- Can show strong logical connections between individuals, events, and evidence

Documentation

- Entire process documented, with each step explainable and justifiable



The Beginning: Incident Alert

System administrator notices strange behavior on a server (slow, hanging...)

Intrusion detection system alerts administrator of suspicious network traffic

Company suddenly loses a lot of sales

Citizen reports criminal activity

- Computer repair center notices child pornography during a computer repair, notifies police

Murder, computer at the scene

Murder, victim has a PDA

Law enforcement: must investigate

Corporate/military: may investigate, depending on severity, other priorities



Crime Scene

Document, document, document!

Photographs depicting the organization of equipment, cabling

Detailed inventory of evidence

Proper handling procedures, turn on, leave off rules for each type of digital device

e.g., for computer:

- Photograph screen, then disconnect all power sources
- Place evidence tape over each drive slot
- Photograph/diagram and label back of computer components with existing connections
- Label all connectors/cable ends to allow reassembly as needed
- If transport is required, package components and transport/store components as fragile cargo



Examples of Digital Evidence

Computers increasingly involved in criminal and corporate investigations

Digital evidence may play a supporting role or be the “smoking gun”

Email

- Harassment or threats
- Blackmail
- Illegal transmission of internal corporate documents



Examples (2)

Meeting points/times for drug deals

Suicide letters

Technical data for bomb making

Image or digital video files (esp., child pornography)

Evidence of inappropriate use of computer resources or attacks

- Use of a machine as a spam email generator
- Use of a machine to distribute illegally copied software



Sources of Digital Evidence

Computers

- Email
- Digital images
- Documents
- Spreadsheets
- Chat logs
- Illegally copied software or other copyrighted material



Digital Evidence on a Disk

Files

- Active
- Deleted
- Fragments

File metadata

Slack space

Swap file

System information

- Registry
- Logs
- Configuration data



More Sources (1)

Wireless telephones

- Numbers called
- Incoming calls
- Voice mail access numbers
- Debit/credit card numbers
- Email addresses
- Call forwarding numbers

PDAs/Smart Phones

- Above, plus contacts, maps, pictures, passwords, documents, ...



More Sources (2)

Landline Telephones/Answering machines

- Incoming/outgoing messages
- Numbers called
- Incoming call info
- Access codes for voice mail systems
- Contact lists

Copiers

- Especially digital copiers, which may store entire copy jobs



More Sources (3)

Video game systems

- Basically computer systems, especially XBox.

GPS devices

- Routes, way-points

Digital cameras

- Photos (obvious) but also video, arbitrary files on storage cards (SD, memory stick, CF, ...)



Preservation of Evidence

Stabilize evidence

Depends on device category, but must keep volatile devices happy

Whenever possible, make copies of original evidence

Write blocking devices and other technology to ensure that evidence is not modified are typically employed

- **Careful! Not all evidence preservation devices work as advertised!**

Original evidence then goes into environmentally-controlled, safe location

“Feeding” of volatile devices continues in storage


Copies of evidence are used for the next phase of investigation

On the Scene Preservation

tick...tick...tick...




“Dear Susan,
It’s not your
fault...”




**Just pull the plug?
Move the mouse for a quick peek?**

**Volatile
computing**



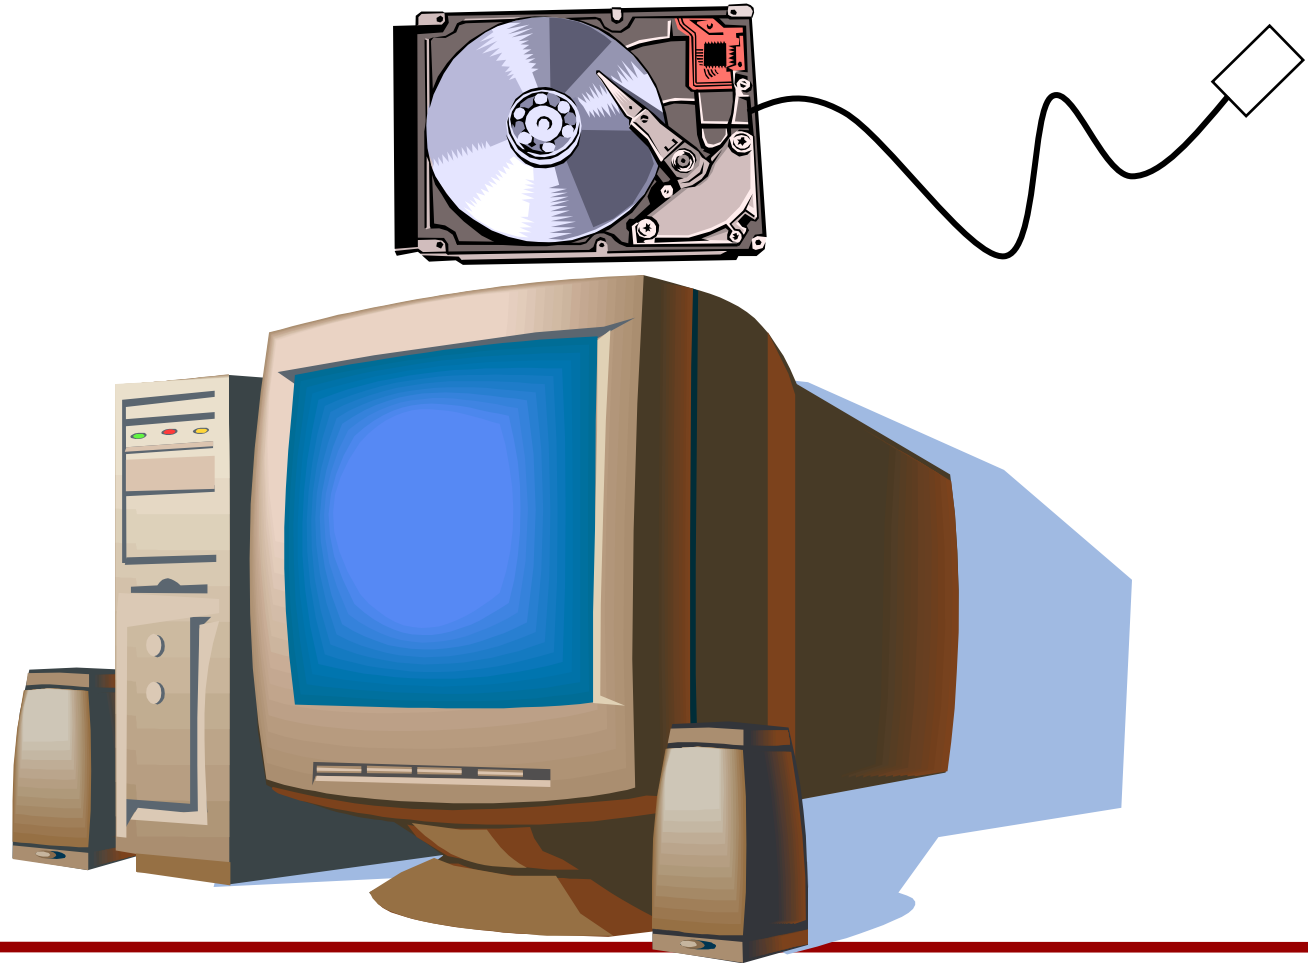
Living room

wireless connection



Basement/closet

Careful Documentation is Crucial



Preservation: Imaging

When making copies of media to be investigated, must prevent accidental modification or destruction of evidence!

- Write blockers: A good plan.

Tools for imaging:

- dd under Linux
- DOS boot floppies
- Proprietary imaging solutions



Drivelock
write blocker



Analysis: Art, Science, Experience

Know where evidence can be found

Understand techniques used to hide or “destroy” digital data

Toolbox of techniques to discover hidden data and recover “destroyed” data

Cope with HUGE quantities of digital data...

Ignore the irrelevant, target the relevant

Thoroughly understand circumstances which may make “evidence” unreliable

- **One example: Creation of new users under Windows 95/98 at the logon**



Traditional Computer: Where's the Evidence?

Undeleted files, expect some names to be incorrect

Deleted files

Windows registry

Print spool files

Hibernation files

Temp files (all those .TMP files in Windows!)

Slack space

Swap files

Browser caches

Alternate or “hidden” partitions

On a variety of removable media (floppies, ZIP, tapes, ...)



Analysis (1)

Using copies of original digital evidence, recover as much evidence as possible

Discovery of deleted files

Discovery of renamed files

Recovery of data blocks for long-deleted files

Discovery of encrypted material

Creation of indices for keyword searches against slack space, swap file, unallocated areas

Use cryptographic hash dictionaries to identify known important/irrelevant files



Analysis (2)

File carving to recover deleted files, file fragments from unallocated space

Discovery of known files using hash dictionaries, to eliminate operating system files, executables for popular application suites, ...

Categorization of evidence

- x JPEG files
- y Word files
- z encrypted ZIP files
- ...

Application of password cracking techniques to open encrypted material

Many of these processes can be automated



Analysis (3)

Creation of a timeline illustrating file creation, modification, deletion dates

For Unix filesystems: inode # “timelines”

Unusual activity will then “pop out” on the timeline

- **Careful! Clock skew, timezone issues, dead CMOS battery...**

Viewing undeleted and recovered data meeting relevant criteria

- e.g., in a child pornography case, look at recovered JPEG/GIF images and any multimedia files
- Probably would not investigate Excel or financial documents

Formulation of hypotheses and the search for additional evidence to justify (or refute) these hypotheses



An Investigative Sampler

Impossible to illustrate many traditional forensics techniques in a short time

Idea: quickly illustrate diversity of available techniques with a few examples

Windows Registry

Swap File

Hibernation File

Recycle Bin

Print Spool Files

Filesystem Internals

File Carving

Slack Space

(similar structures on Linux, Mac OS X, etc.)



FTK Screenshots: Thumbnail View

AccessData FTK version 1.62.1 build 06.07.27 -- C:\Research\tutorials\sacsac2006\DARYL_EVIDENCE\FTK\DarylPoppins\

File Edit View Tools Help

Overview Explore Graphics E-Mail Search Bookmark

11653 Listed 0 Checked Total 6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\propviews[88].gif

File Name	Full Path	Recycle ...	Ext	File Type	Category	Subject	Cr Date	Mod Date	Acc Date	L-Size	P-Size	Childr...	Desc...	Enc	D...	Re...	Cr
propviews[84].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:32:25 ...	7/21/2003 12:32:25 ...	7/21/2003 12:32:25 ...	43	43	0	0	0	Y		
propviews[84].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:34:14 ...	7/21/2003 12:34:14 ...	7/21/2003 12:34:14 ...	43	43	0	0	0	Y		
propviews[84].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:44:20 ...	7/21/2003 12:44:21 ...	7/21/2003 12:44:21 ...	43	43	0	0	0	Y		
propviews[84].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:52:09 ...	7/21/2003 12:52:09 ...	7/21/2003 12:52:09 ...	43	43	0	0	0	Y		
propviews[85].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:36:40 ...	7/21/2003 12:36:40 ...	7/21/2003 12:36:40 ...	43	43	0	0	0	Y		
propviews[85].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:36:53 ...	7/21/2003 12:36:54 ...	7/21/2003 12:36:54 ...	43	43	0	0	0	Y		
propviews[85].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:47:27 ...	7/21/2003 12:47:27 ...	7/21/2003 12:47:27 ...	43	43	0	0	0	Y		
propviews[85].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:55:42 ...	7/21/2003 12:55:42 ...	7/21/2003 12:55:42 ...	43	43	0	0	0	Y		
propviews[86].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:38:41 ...	7/21/2003 12:38:41 ...	7/21/2003 12:38:41 ...	43	43	0	0	0	Y		
propviews[86].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:40:19 ...	7/21/2003 12:40:19 ...	7/21/2003 12:40:35 ...	43	43	0	0	0	Y		
propviews[86].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:51:06 ...	7/21/2003 12:51:06 ...	7/21/2003 12:51:06 ...	43	43	0	0	0	Y		
propviews[86].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:58:24 ...	7/21/2003 12:58:24 ...	7/21/2003 12:58:32 ...	43	43	0	0	0	Y		
propviews[87].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:51:38 ...	7/21/2003 12:51:38 ...	7/21/2003 12:51:38 ...	43	43	0	0	0	Y		
propviews[87].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 1:03:31 AM	7/21/2003 1:03:31 AM	7/21/2003 1:03:31 AM	43	43	0	0	0	Y		
propviews[87].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:39:14 ...	7/21/2003 12:39:14 ...	7/21/2003 12:39:14 ...	43	43	0	0	0	Y		
propviews[87].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:42:14 ...	7/21/2003 12:42:14 ...	7/21/2003 12:42:14 ...	43	43	0	0	0	Y		
propviews[88].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 1:06:23 AM	7/21/2003 1:06:23 AM	7/21/2003 1:06:23 AM	43	43	0	0	0	Y		
propviews[88].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:53:54 ...	7/21/2003 12:53:54 ...	7/21/2003 12:53:54 ...	43	43	0	0	0	Y		
propviews[88].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:42:31 ...	7/21/2003 12:42:31 ...	7/21/2003 12:42:31 ...	43	43	0	0	0	Y		
propviews[88].gif	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		gif	GIF File	Graphic		7/21/2003 12:44:40 ...	7/21/2003 12:44:40 ...	7/21/2003 12:44:40 ...	43	43	0	0	0	Y		
propviews[89].nil	6gb-fixed-chkdsk\NONAME-NTFS\{orphan}\pro...		nil	GIF File	Graphic		7/21/2003 12:47:09 ...	7/21/2003 12:47:09 ...	7/21/2003 1:02:16 AM	43	43	0	0	0	Y		



Windows Registry

Can be a forensics goldmine

Lots of information, fairly difficult to “clean”

Username

Internet history

Program installation information

Recently accessed files

USB device history



More Registry

Other useful info obtainable from the registry:

- CPU type
- Network interface information
 - IP addresses, default gateway, DHCP configuration, ...
- Installed software
- Installed hardware

Registry information “gotchas”

- redundant, undocumented information
- profile cloning on older versions of Windows (95/98)
 - (e.g., typed URLs, browser history, My Documents, ...)



File Systems

Quick overview



File Systems

Data

- Files
- Directories

Metadata

- Time stamps (modify, access, create/change, delete)
- Owner
- Security properties

Structures

- Superblock/Master File Table/File Access Table
- inodes/clusters
- data



File Systems (2)

More sophisticated data recovery requires deep knowledge of filesystem internals

Structures that manage filesystem metadata

Disk layout

File deletion issues

Many important filesystems

- DOS / Windows: FAT, FAT16, FAT32, NTFS
- Unix: ext2, ext3, Reiser, JFS, ... more
- Mac: MFS, HFS, HFS+



File Systems: ext2 and ext3

Efficient file system, supports

- indirect blocks (double and triple indirection)
- symbolic links
- sparse files

Has MAC times, but no file creation time

ext3 = ext2 + journaling for faster crash recovery and system boot file check

Forensic artifacts from file deletion:

- ext2: content preserved, connection to name lost
- ext3: connection to content lost, metadata preserved



File Deletion: Linux

ext2 file deletion

- Adjust previous directory entry length to obscure deleted record
- No reorganization to make space in directories
- “first fit” for new directory entries, based on real name length
- Directory entry’s inode # is cleared

ext3 file deletion

- Same as for ext2, but...
 - inode is wiped on file deletion, so block numbers are lost
 - Major anti-forensics issue!
 - But directory entry’s inode # isn’t cleared...
-



File Systems: FAT

FAT12, FAT16, FAT32

- different size of addressable cluster

Common format for floppy disks (remember those?)

Limited time/date information for FAT files

- Last write date/time is always available
- Creation date/time is optional and may not be available
- Last access DATE ONLY is optional and may not be available

Short file names (8.3) on FAT12 and FAT16

No security features

Long names for FAT32



FAT: Short Filename Storage

- "foo.bar" -> "FOO BAR"
- "FOO.BAR" -> "FOO BAR"
- "Foo.Bar" -> "FOO BAR"
- "foo" -> "FOO "
- "foo." -> "FOO "
- "PICKLE.A" -> "PICKLE A "
- "prettybg.big" -> "PRETTYBGBIG"

Note case is not significant

“.” between primary filename and extension is implied (not actually stored)

Further, everything is space-padded



FAT: More Dir Entry Details

Date format:

- Bits 0–4: Day of month, valid value range 1-31 inclusive.
- Bits 5–8: Month of year, 1 = January, valid value range 1–12 inclusive.
- Bits 9–15: Count of years from 1980, valid value range 0–127 inclusive (1980–2107).

Time Format:

- A FAT directory entry time stamp is a 16-bit field that has a granularity of 2 seconds
- Bits 0–4: 2-second count, valid value range 0–29 inclusive (0 – 58 seconds).
- Bits 5–10: Minutes, valid value range 0–59 inclusive
- Bits 11–15: Hours, valid value range 0–23 inclusive



FAT: Long Filenames

Summary: a kludge to add support without changing short-name handling

Up to 255 characters in pathname component

Total pathname no longer than 260

More supported characters

Leading/trailing spaces ignored

Internal spaces allowed

Leading/embedded “.” allowed

Trailing “.” are ignored

Stored case-sensitive

Processed case-insensitive (for compatibility)

File created with short name (uses “~1”, “~2”, etc. suffix)



File Systems: NTFS

Master File table grows, never shrinks (artifacts!)

B-tree algorithm used for file tree

- re-“balances” file system tree when tree changes
- creating or deleting a file can cause entire tree to change and can overwrite nodes that were marked as free but still had information in them
- can destroy artifacts!

lots of attributes on files, can be confusing (e.g., which access time is the “official” one to use)

- most useful attributes are MAC times



File Systems: Partitions

Physical disk divided into logical partitions

Logical partitions may not be mounted or may be in a format the running O/S does not recognize (e.g., dual boot system)

Formats:

- DOS (most common)
- Apple
- Solaris
- BSD
- RAID (can cause difficulties for investigators if disk slices have to be reconstructed manually)



File System Forensic Artifacts

Active files

- contents (data blocks)
- metadata (owner, MAC times)
- permissions (ACLs)
- who is using it now (***not*** in a static analysis)

Deleted files

- full contents (sometimes, depends on usage)
- partial contents (via carving)
- metadata (sometimes, depends on O/S)
 - deletion times



File Deletion: Windows

FAT file deletion

- Directory entry has first character changed to 0xE5
- Directory entry contains first cluster number (index into FAT); this isn't lost when file is deleted
- Other FAT entries for file are cleared

NTFS file deletion

- IN_USE flag on MFT entry for file is cleared
- Parent directory entry is removed and directory is re-sorted
- Data clusters marked as unallocated
- Filename is likely to be lost, but since MFT entry isn't destroyed, file data may be recoverable
- Dates aren't lost
- Caveat: NTFS reuses MFT entries before creating new ones, so recoverable deleted files are probably recently deleted ones



File Rename, Move

When a file is renamed under Windows, old directory entry is deleted and new one created

Starting cluster is the same for each

Establishing that a user moved or renamed a file can provide evidence that the user knew of the file's existence



Useful Files with Forensic Content



Windows Recycle Bin

Indirect file deletion facility

Mimics functionality of a trashcan

- Place “garbage” into the can
- You can change your mind about the “garbage” and remove it, until...
- ...trash is emptied, then it’s “gone”

Files are moved into a special directory

Deleted only when user empties



Windows Shortcut Files

In Desktop, Recent, etc. directories

*.lnk files

Give information about configuration of desktop

Existence of desktop shortcuts (even if the shortcut files are deleted) can...

...establish that user knew of the existence of the files

...establish that user organized files

e.g., can be used to dismiss claims that child pornography or illegal copies of software were “accidentally” downloaded in a bulk download operation



Windows Swap Files

Supports Windows virtual memory system

Contains swapped out pages corresponding to executing processes

NT, Win2000, XP

- Generally, `c:\pagefile.sys`
- Hidden file

95/98

- `c:\windows\win386.swp`
- Hidden file



Windows Swap File: Overview

Potentially, contains a lot of junk

File carving or keyword searches against the raw disk will yield a superset of the information in the swap file (obviously)

May be useful to target swap file directly, particularly on large drives

Careful!

Keyword matches against the swap file **DO NOT necessarily** mean that the corresponding strings were in pages swapped out during the last boot!

When the swap file is created, the “underlying” blocks aren’t cleaned

As the swap file is reused, not all blocks are cleaned

- Swap file can create a “jail”, where e.g. deleted file data from the browser cache end up “trapped” in the set of blocks allocated to the swap file

Blocks may not be overwritten even during months of use!



Hibernation Files

Memory image of XP box, created at shutdown

Allows fast restart

Hibernation file locked during OS execution

Approximately the size of physical RAM (e.g., 2GB RAM == ~2GB hibernation file)

Potentially much more interesting than swap file, since it allows the last “on” state of the machine to be recreated



Hibernation (2)

Can search hibernation file for interesting strings, including URLs, passwords, etc.

First block of file is zero-filled after boot, so you get one chance to “boot” the machine again, unless you have a backup of the hibernation file

Remainder of hibernation file remains unchanged until another hibernation event occurs...

Means that you may be able to recover interesting information that is quite old (reincarnation attack strikes back)



Windows Print Spool Files

*.spl, *.shd files

.shd file contains information about the file being printed

.spl file contains info to render the contents of the file to be printed

Presence of .shd files can be used in a similar fashion as for shortcut files...

...shows knowledge of existence of files and a deliberate attempt to access (print) the contents of the file



Analysis: Evidence Correlation

- Chat logs for IRC channel catering to trading of illegally copied software
- File creation dates for illegal software close to those of the chat session
- Bulk downloads of illegal images followed by categorization of images
- Incriminating categories (e.g., directories)?
- Correlation is still largely a human task



Analysis: Challenges (1)

Digital evidence: incomplete view of communication

Example:

- Digital communication event between two human beings
- Primary method: EMAIL
- Hundreds or thousands of keystrokes and mouse clicks, which were probably **not** captured
- Draft copies of email which may not represent the actual message that was sent
- Fragments of email in browser cache (for web-based email)
- Attachments?
- Secondary communication streams during event?
- Messenger programs (e.g., “I’m sending you that suicide letter I wrote as my creative writing project...”)



Analysis: Challenges (2)

Interactions with other computers

Internet makes investigation much more difficult

Use of encryption, steganography

“Secure” deletion

- Luckily (?) some secure deletion software is horribly broken

Operating systems features!

- e.g., ext3 filesystem in Linux
- Secure recycle bin in Mac OS X

Criminals are getting smarter, many current investigative techniques will need to be improved



Reporting

Case reports must include detailed explanations of every step in the investigative process

Detail must be sufficient to recreate the entire process

- ...
- A keyword search on “heroin” revealed a deleted email message with an attachment as well as a number of other email messages in which an alias was used by the defendant
- The attachment on the matching email file was an encrypted ZIP archive named “credits.zip”
- Attempts to crack the ZIP password using the Password Recovery Toolkit failed to reveal the password, so a number of aliases used by the suspect in the emails were tried as passwords
- “trainspotter” was discovered to be the ZIP password
- Located inside the ZIP file was a text file with a number of credit card numbers, none of which were found to belong to the defendant



More Sophisticated “Dead” Analysis

File carving = rebuilding file even when metadata is not available

Better auditing of investigative process

Better (automated) correlation of evidence

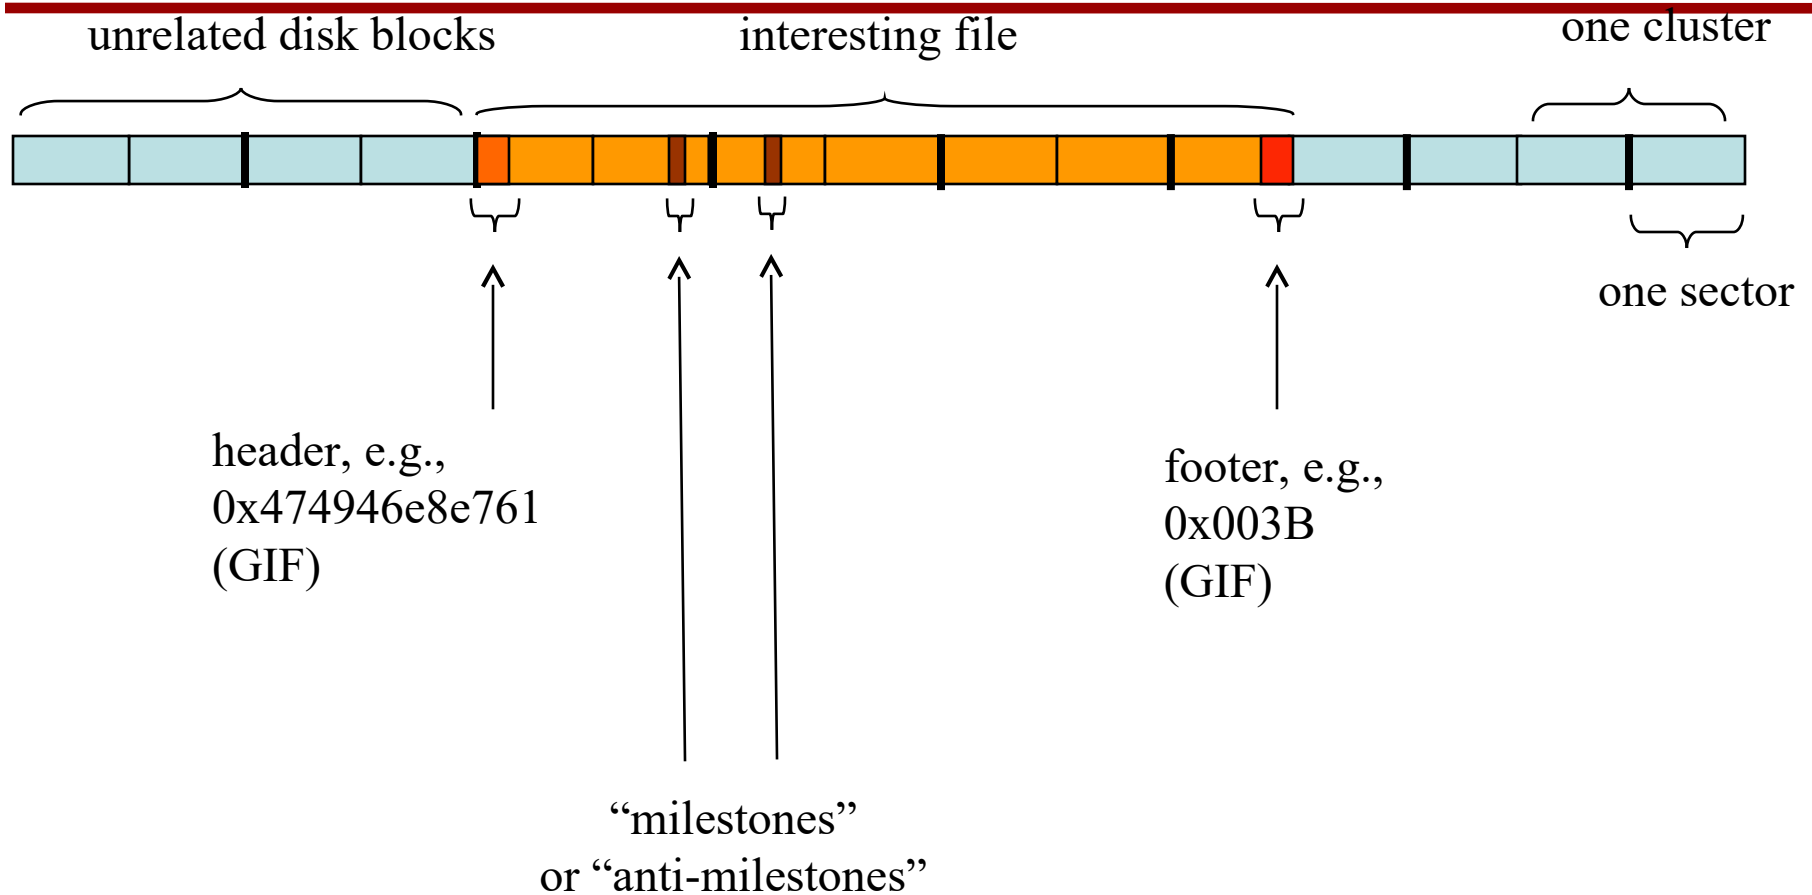
Better handling of multimedia

Distributed digital forensics


Massively threaded digital forensics tools

- e.g., GPUs, multicore CPUs

File Carving: Basic Idea

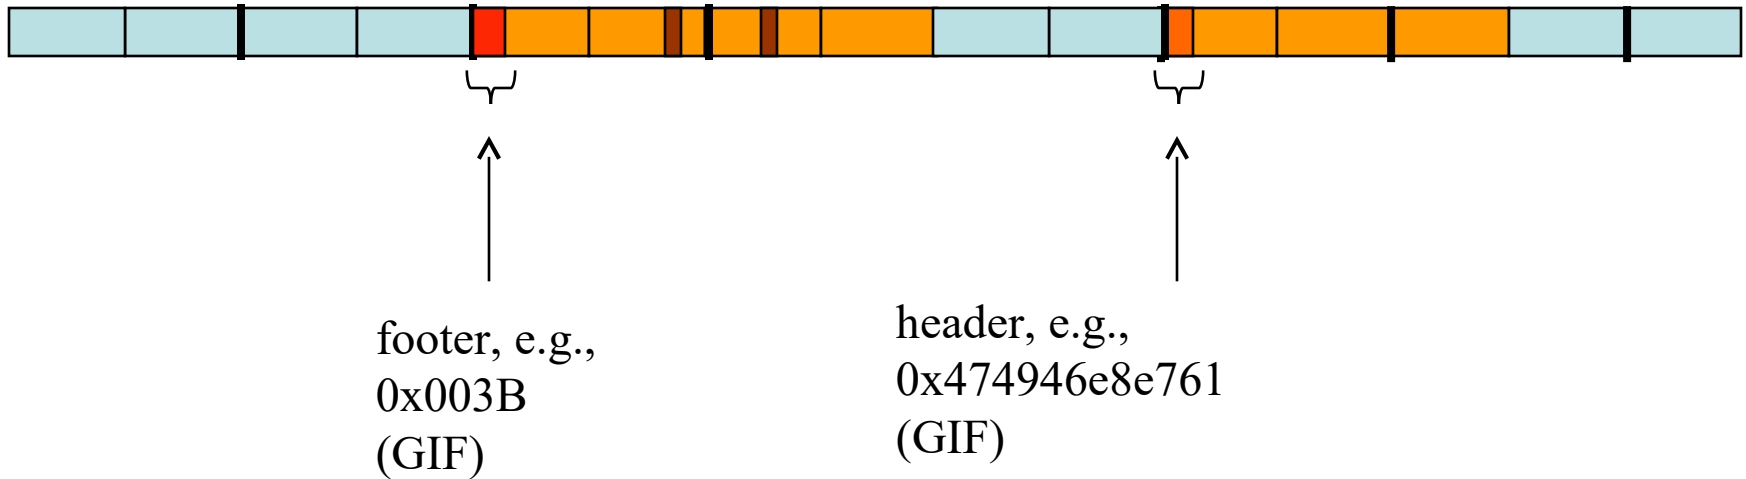


File Carving: Fragmentation






File Carving: Fragmentation

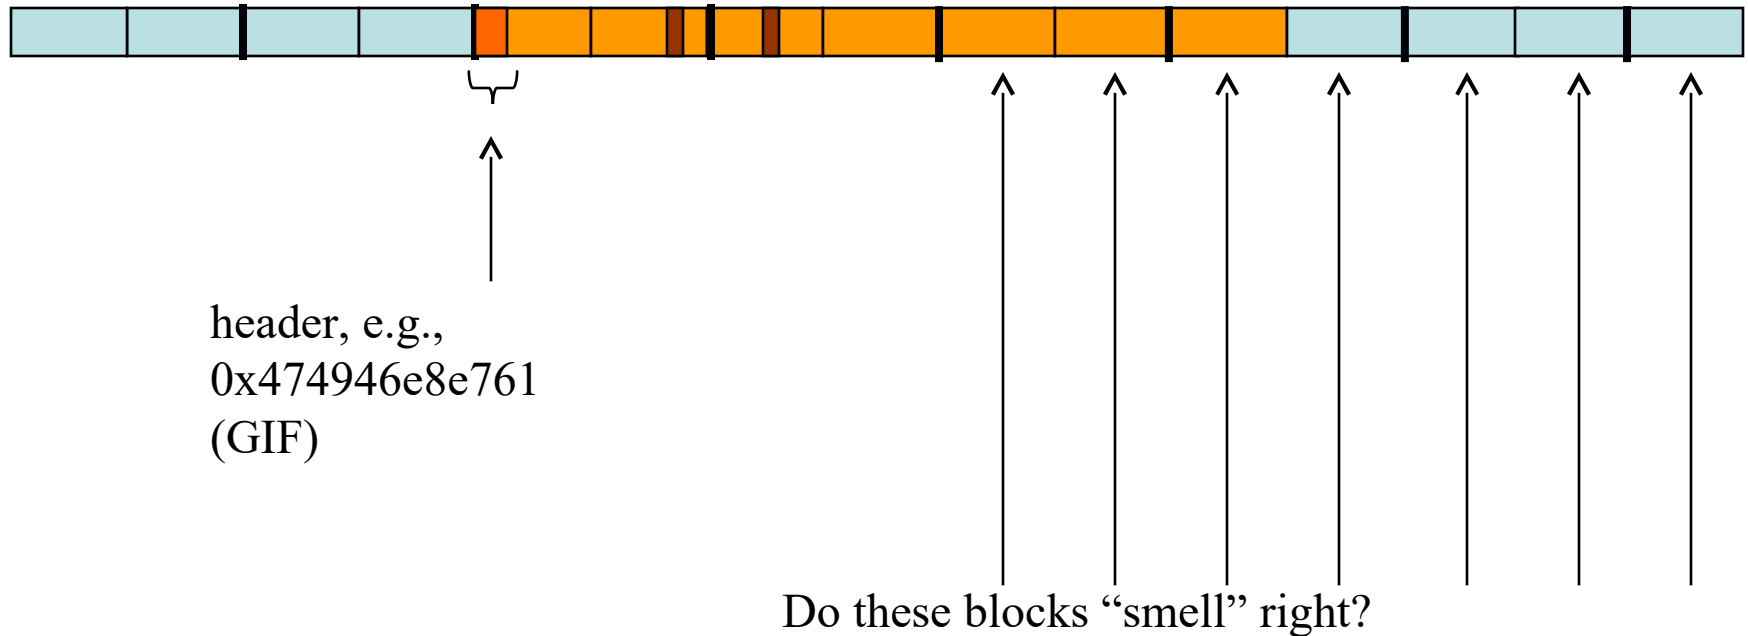




File Carving: Damaged Files



File Carving: Block Sniffing



- N-gram analysis
- entropy tests
- deep analysis



But Evidence is Also...

Improving performance and sophistication of “dead” forensics is important, but evidence is also...

“In” the network

In RAM

On machine-critical machines

- Can’t turn off without severe disruption
- Can’t turn them ALL off just to see!

On huge storage devices

- 1TB server: image entire machine and drag it back to the lab to see if it’s interesting?
- 10TB?



Simple Network Forensics

Obtain another piece of the puzzle

Find information on “what happened” by looking in the network packet flow

Information can be used to:

- Reconstruct sessions (e.g., web, ftp, telnet, IM)
- Find files (downloaded or accessed through network drives)
- Find passwords
- Identify remote machines



Constraints

Legal

- While there is a wealth of information on the network, there are **MANY** legal constraints relating to wire-tapping, e.g.,
 - Computer Fraud and Abuse Act (18 U.S.C. § 1030)
 - Electronic Communications Privacy Act ("ECPA"), 18 U.S.C. § 2703 et seq)
 - "wire communication" (18 U.S.C. § 2510)
 - plus state laws
- May depend on what information you collect, whether it is part of the normal practices, whether there is any “reasonable expectation to privacy,” etc.
- The laws can be subtle...
- Consult an expert first and have a policy defined ahead of time!



Constraints

Technical

- tapping the right line
 - switched vs. flat networks
- determining proper IP addresses
- IP addresses may change over time
- corroborating evidence with:
 - log files
 - evidence obtained from traditional forensic evaluation
 - evidence obtained from live forensic evaluation
- encrypted data



Typical Scenario

“Dead” forensics information incomplete

- discovered to be incomplete
- predicted to be incomplete

Non-local attacker or local user using network in inappropriate fashion

Generally, another event triggers network investigation

Company documents apparently stolen

Denial of service attack

Suspected use of unauthorized use of file sharing software

“Cyberstalking” or threatening email



Information Available

Summary information (router flow logs)

- Routers generally provide this information
- Includes basic connection information
 - source and destination IP address and ports
 - connection duration
 - number of packets sent
- No content! Can only surmise what was sent
- Can establish that connections between machines were established
- Can corroborate data from log files (e.g., ssh'ing from one machine to another to another within a network)
- Unusual ports (rootkits? botnet?)
- Unusual activity (spam generator?)



Information Available (2)

Complete information (packet dumps)

- from programs like Ethereal/Wireshark, snort, tcpdump
- on an active net, can generate a **LOT** of data
- can provide filter options so programs only capture certain traffic (by IP, port, protocol)
- includes full content—can reconstruct what happened (maybe)
- reconstruct sessions
- reconstruct transmitted files
- retrieve typed passwords
- identify which resources are involved in attack
- BUT no easy way to decrypt encrypted traffic



Information Available (3)

Port scans (nmap, etc.)

- Identifies machines on your network
 - Often can identify operating system, printer type, etc., **without** needing account on the machine
 - “OS fingerprinting”
- Identifies ports open on those machines
 - Backdoors, unauthorized servers, ...
- Identifies suspicious situation (infected machine, rogue computer, etc.)
- nmap: lots of options



Analysis

Does not exist in a vacuum

Link information in analysis to network and host log files

- who was on the network
- who was at the keyboard
- what files are on the disk and where

Look up the other sites (who are they, where are they, what's the connection)

Otherwise, network traces can be overwhelming

Potentially huge amounts of data

Limited automation!



Normal ICMP Traffic (tcpdump)

Pings

```
IP BOUDIN.mshome.net > www.google.com: icmp 40: echo request seq 6400
IP www.google.com > BOUDIN.mshome.net: icmp 40: echo reply seq 6400
IP BOUDIN.mshome.net > www.google.com: icmp 40: echo request seq 6656
IP www.google.com > BOUDIN.mshome.net: icmp 40: echo reply seq 6656
IP BOUDIN.mshome.net > www.google.com: icmp 40: echo request seq 6912
IP www.google.com > BOUDIN.mshome.net: icmp 40: echo reply seq 6912
IP BOUDIN.mshome.net > www.google.com: icmp 40: echo request seq 7168
IP www.google.com > BOUDIN.mshome.net: icmp 40: echo reply seq 7168
```

Host unreachable

```
xyz.com > boudin.cs.uno.edu: icmp: host blarg.xyz.com unreachable
```

Port unreachable

```
xyz.com > boudin.cs.uno.edu: icmp: blarg.xyz.com port 7777 unreachable
```



Fragmentation Visualization

Fragmentation can be seen by tcpdump

```
whatever.com > me.com: icmp: echo request (frag 5000:1400@0+)  
whatever.com > me.com: (frag 5000:1000@1400)
```

ID

size

offset

more frags flag

Note that 2nd frag
isn't identifiable as ICMP
echo request...



nmap 137.30.120.*

```
Starting Nmap 4.11 ( http://www.insecure.org/nmap )  
  at 2006-10-24 19:32
```

```
Interesting ports on 137.30.120.1:
```

```
Not shown: 1679 closed ports
```

```
PORT      STATE SERVICE
```

```
23/tcp    open  telnet
```

```
MAC Address: 00:0D:ED:41:A8:40 (Cisco Systems)
```

```
All 1680 scanned ports on 137.30.120.3 are closed
```

```
MAC Address: 00:0F:8F:34:7E:C2 (Cisco Systems)
```

```
All 1680 scanned ports on 137.30.120.4 are closed
```

```
MAC Address: 00:13:C3:13:B4:41 (Cisco Systems)
```

```
All 1680 scanned ports on 137.30.120.5 are closed
```

```
MAC Address: 00:0F:90:84:13:41 (Cisco Systems)
```

```
...
```



Wireshark (aka Ethereal)

Packet listing

No.	Time	Source	Destination	Src Port	Dest Port	Protocol	Info
100	20.143472	64.233.167.99	137.30.123.234	2157	80	HTTP	[TCP out-of-order] HTTP/1.1 200 OK (text/html)
101	28.029940	137.30.123.234	64.233.167.99	2159	http	HTTP	GET / HTTP/1.1
95	21.764143	216.239.51.99	137.30.123.234	http	2161	HTTP	[TCP Previous segment lost] Continuation or non-HTTP traffic
85	20.811074	137.30.123.234	216.239.51.99	2161	http	HTTP	GET /groups?q=rhino.exe&hl=en&lr=&ie=UTF-8&oe=UTF-8&selm=MPG.159359f528ff3d2798aaca%40news.ntl
76	14.500008	216.239.51.99	137.30.123.234	http	2161	HTTP	[TCP Previous segment lost] Continuation or non-HTTP traffic
74	14.499702	216.239.51.99	137.30.123.234	http	2161	HTTP	[TCP Previous segment lost] Continuation or non-HTTP traffic
73	14.436408	137.30.123.234	216.239.51.99	2161	http	HTTP	GET /intl/en_ALL/images/groups_res.gif HTTP/1.1
68	14.314018	216.239.51.99	137.30.123.234	http	2161	HTTP	[TCP Previous segment lost] Continuation or non-HTTP traffic
59	13.922479	137.30.123.234	216.239.51.99	2161	http	HTTP	GET /groups?q=rhino.exe&hl=en&lr=&ie=UTF-8&oe=UTF-8&sa=N&tab=wg HTTP/1.1
54	8.878986	64.233.167.99	137.30.123.234	http	2159	HTTP	[TCP Retransmission] HTTP/1.1 200 OK (GIF89a)
50	8.655695	64.233.167.99	137.30.123.234	http	2157	HTTP	[TCP Previous segment lost] Continuation or non-HTTP traffic
49	8.581627	137.30.123.234	64.233.167.99	2157	http	HTTP	GET /nav_next.gif HTTP/1.1
47	8.326670	64.233.167.99	137.30.123.234	http	2157	HTTP	HTTP/1.1 200 OK (GIF89a)
46	8.243665	137.30.123.234	64.233.167.99	2157	http	HTTP	GET /nav_page.gif HTTP/1.1
40	7.987860	64.233.167.99	137.30.123.234	http	2157	HTTP	HTTP/1.1 200 OK (GIF89a)
38	7.913810	137.30.123.234	64.233.167.99	2157	http	HTTP	GET /nav_current.gif HTTP/1.1
35	7.743128	137.30.123.234	64.233.167.99	2159	http	HTTP	GET /images/logo_sm.gif HTTP/1.1
31	7.659537	64.233.167.99	137.30.123.234	http	2157	HTTP	HTTP/1.1 200 OK (GIF89a)
30	7.583199	137.30.123.234	64.233.167.99	2157	http	HTTP	GET /nav_first.gif HTTP/1.1
10	7.043566	137.30.123.234	64.233.167.99	2157	http	HTTP	GET /search?hl=en&ie=UTF-8&q=rhino.exe HTTP/1.1
8	0.178056	64.233.167.99	137.30.123.234	http	2157	HTTP	[TCP out-of-order] HTTP/1.1 200 OK (text/html)
4	0.070991	137.30.123.234	64.233.167.99	2157	http	HTTP	GET / HTTP/1.1

Detailed packet data at various protocol levels

Frame 49 (443 bytes on wire, 443 bytes captured)

- Ethernet II, Src: Applecom_cc:57:92 (00:03:93:cc:57:92), Dst: Cisco_41:a8:40 (00:0d:ed:41:a8:40)
- Internet Protocol, Src: 137.30.123.234 (137.30.123.234), Dst: 64.233.167.99 (64.233.167.99)
- Transmission Control Protocol, Src Port: 2157 (2157), Dst Port: http (80), Seq: 2012, Ack: 20812, Len: 389**
 - Source port: 2157 (2157)
 - Destination port: http (80)
 - Sequence number: 2012 (relative sequence number)
 - [Next sequence number: 2401 (relative sequence number)]
 - Acknowledgement number: 20812 (relative ack number)
 - Header length: 20 bytes
 - Flags: 0x0018 (PSH, ACK)
 - Window size: 63662
 - Checksum: 0xeef4 [incorrect, should be 0x28d3 (maybe caused by checksum offloading?)]**
- Hypertext Transfer Protocol**
 - GET /nav_next.gif HTTP/1.1\r\n**
 - Accept: /*/*\r\n
 - \r\n
 - Accept-Language: en-us\r\n
 - \r\n
 - User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; windows NT 5.1; .NET CLR 1.1.4322)\r\n
 - Host: www.google.com\r\n
 - Connection: Keep-Alive\r\n
 - Cookie: PREF=ID=269e53562e69c3c5:TM=1082665356:LM=1083170916:TB=3:5=53_bmTawajyv1FX0\r\n\r\n

Raw data

0030 f8 ae ee f4 00 00 47 45 54 20 2f 6e 61 76 5f 6eGE T /nav_n
0040 65 78 74 2e 67 69 66 20 48 54 54 50 2f 31 2e 31 ext.gif HTTP/1.1
0050 0d 0a 41 63 63 65 70 74 3a 20 2a 2f 2a 0d 0a 2d ..Accept: /*/*..
0060 2d 2d 2d 2d 2d 2d 3a 20 2d 2d 2d 2d 3a 2d 2d 2d -----:-----
0070 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d -----
0080 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d 2d -----

Hypertext Transfer Protocol (http), 389 bytes | P: 295 D: 295 M: 0



Wireshark: Following a TCP Stream

rhino.log - Wireshark

Filter: Expression... Clear Apply

No.	Time	Source	Destination	Src Port	Dest Port	Protocol	Info
1737	208.081951	137.30.120.40	137.30.122.253	ftp	1658	FTP	Response: 226-WARNING! 321 bare linefeeds received in ASCII mode
1653	207.979567	137.30.120.40	137.30.122.253	ftp	1658	FTP	Response: 150 opening ASCII mode data connection for rhino3.jpg.
1649	207.947603	137.30.122.253	137.30.120.40	1658	ftp	FTP	Request: STOR rhino3.jpg
1648	207.946973	137.30.120.40	137.30.122.253	ftp	1658	FTP	Response: 200 PORT command successful.
1647	207.945618	137.30.122.253	137.30.120.40	1658	ftp	FTP	Request: PORT 137,30,122,253,6,124
1631	200.357806	137.30.120.40	137.30.122.253	ftp	1658	FTP	Response: 230 User gnome logged in.
1629	200.280951	137.30.122.253	137.30.120.40	1658	ftp	FTP	Request: PASS gnome123
1627	198.529854	137.30.120.40	137.30.122.253	Tcp	1658	Tcp	Password required for gnome.
1625	198.525443	137.30.122.253	137.30.120.40	1658	ftp	FTP	gnome
1623	195.462335	137.30.120.40	137.30.122.253	ftp	1658	FTP	cook FTP server ready.
1616	194.432107	137.30.120.40	137.30.122.253	ftp	1658	FTP	-Total traffic for this session was 66042 bytes in 1 transfers.
1615	194.427484	137.30.120.40	137.30.122.253	ftp	1658	FTP	-You have transferred 65703 bytes in 1 files.
1614	194.426879	137.30.122.253	137.30.120.40	1655	ftp	FTP	Transfer complete.
1612	189.221711	137.30.120.40	137.30.122.253	ftp	1658	FTP	Opening BINARY mode data connection for rhino1.jpg.
1550	189.033465	137.30.120.40	137.30.122.253	ftp	1658	FTP	rhino1.jpg
1546	188.996081	137.30.122.253	137.30.120.40	1658	ftp	FTP	PORT command successful.
1545	188.995519	137.30.120.40	137.30.122.253	ftp	1658	FTP	137,30,122,253,6,121
1544	188.994914	137.30.122.253	137.30.120.40	1655	ftp	FTP	Type set to I.
1541	185.602818	137.30.120.40	137.30.122.253	ftp	1658	FTP	User gnome logged in.
1540	185.602553	137.30.122.253	137.30.120.40	1658	ftp	FTP	
1538	184.748946	137.30.120.40	137.30.122.253	ftp	1658	FTP	

Frame 1627 (88 bytes on wire, 88 bytes captured)

Ethernet II, Src: sunMicro_f0:13:96 (08:00:20:f0:13:96), Dst: AppleCom_03:57:92 (00:09:93:cc:57:92)

Internet Protocol, Src: 137.30.120.40 (137.30.120.40), Dst: 137.30.122.253 (137.30.122.253)

Transmission Control Protocol, Src Port: ftp (21), Dst Port: 1658 (1658), Seq: 29, Ack: 13, Len: 34

Source port: ftp (21)

Destination port: 1658 (1658)

Sequence number: 29 (relative sequence number)

[Next sequence number: 63 (relative sequence number)]

Acknowledgement number: 13 (relative ack number)

Header length: 20 bytes

Flags: 0x0018 (PSH, ACK)

window size: 49640

Checksum: 0x19a3 [correct]

File Transfer Protocol (FTP)

0000 00 03 93 cc 57 92 08 00 20 f0 13 96 08 00 45 00W... ..E.
0010 00 4a 80 08 40 00 3c 06 b9 43 89 1e 78 28 89 1e .J..@.<. .C..x(..
0020 7a fd 00 15 06 7a 0e 15 fe 76 51 1a 65 14 50 18 z....Z...vQ.e.P.
0030 c1 e8 19 a3 00 00 33 33 31 20 50 61 73 73 77 6f33 1 Passwo
0040 72 64 20 72 65 71 75 69 72 65 64 20 66 6f 72 20 rd requi red For
0050 67 62 6f 6d 65 20 0d 0a 67 62 6f 6d 65 20 0d 0a gnome

File: "C:\class\4621\au06\rhino.log" 3113 KB 00:12:30 P: 6557 D: 6557 M: 0



Wireshark: FTP Data Stream

(Untitled) - Wireshark

File Edit View Go Capture Analyze Statistics Help

Filter: (ip.addr eq 137.30.123.151 and ip.addr eq 137.30.120.40) and (t... Expression... Clear Apply

No.	Time	Source	Destination	Src Port	Dest Port	Protocol	Info
421	38.321288	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
422	38.321381	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
423	38.321474	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
424	38.321568	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
425	38.321646	137.30.120.40	137.30.123.151	5002	ftp-data	TCP	ftp-data > 5002 [ACK] Seq=1 Ack=77216 win=45260 Len=0
426	38.321718	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
427	38.321811	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
428	38.321903	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
429	38.321997	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
430	38.322090	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
431	38.322183	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
432	38.322275	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
433	38.322402	137.30.120.40	137.30.123.151	5002	ftp-data	TCP	ftp-data > 5002 [ACK] Seq=1 Ack=88896 win=48180 Len=0
434	38.322417	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
435	38.322477	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
436	38.322571	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
437	38.322663	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
438	38.322755	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
439	38.322848	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
440	38.322942	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1460 bytes
441	38.323036	137.30.123.151	137.30.120.40	5002	ftp-data	FTP-DATA	FTP Data: 1262 bytes
442	38.323171	137.30.120.40	137.30.123.151	5002	ftp-data	TCP	ftp-data > 5002 [ACK] Seq=1 Ack=100576 win=43800 Len=0

Frame 421 (1514 bytes on wire, 1514 bytes captured)

- Ethernet II, Src: Ibm_10:93:ac (00:0d:60:10:93:ac), Dst: SunMicro_09:4b:61 (00:03:ba:09:4b:61)
- Internet Protocol, Src: 137.30.123.151 (137.30.123.151), Dst: 137.30.120.40 (137.30.120.40)
- Transmission Control Protocol, Src Port: 5002 (5002), Dst Port: ftp-data (20), Seq: 74296, Ack: 1, Len: 1460
- FTP Data

```
0000 00 03 ba 09 4b 61 00 0d 60 10 93 ac 08 00 45 00  ....Ka..y....E.
0010 05 dc 04 a6 40 00 80 06 ea 79 89 1e 7b 97 89 1e  ....@...y. {...
0020 78 28 13 8a 00 14 26 45 96 b1 47 2c 42 73 50 10  x(...&E...G,BSP.
0030 ff ff 89 48 00 00 d5 ae 78 80 08 e0 76 48 d8 4c  ...H....x...vH.L
0040 ac 30 08 ed b5 01 c7 07 de a7 e6 01 87 99 13 05  .0.....
0050 db 01 02 c8 00 03 e3 d6 24 2d 14 d2 c0 c3 76 c0  .
```

File: "c:\temp\etherXXXX75FDIT" 248 KB 00:00:43 | P: 490 D: 90 M: 0 Drops: 0



Wireshark: HTTP Session

Stream Content

```
GET /icons/back.gif HTTP/1.1
Accept: */*
-----:-----
Accept-Language: en-us
-----:-----
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; windows NT 5.1; .NET CLR 1.1.4322)
Host: www.cs.uno.edu
Connection: Keep-Alive

HTTP/1.1 200 OK
Date: wed, 28 Apr 2004 21:07:26 GMT
Server: Apache/1.3.29 (Unix)
Last-Modified: Thu, 22 Feb 1996 11:45:53 GMT
ETag: "b88a-d8-312c5771"
Accept-Ranges: bytes
Content-Length: 216
Keep-Alive: timeout=15, max=100
Connection: Keep-Alive
Content-Type: image/gif

GIF89a.....fff333.....!.NThis art is in the public domain. Kevin
Hughes, kevinh@eit.com, september
1995.!......K....#.j.3S.....i.7N.v..t..7*E.nXm/...5FP...x....l.
(N. ....: @..V...U.....;GET /~gnome/rhino5.gif HTTP/1.1
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, application/vnd.ms-excel,
application/vnd.ms-powerpoint, application/msword, application/x-shockwave-flash, */*
-----:-----
Accept-Language: en-us
-----:-----
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; windows NT 5.1; .NET CLR 1.1.4322)
Host: www.cs.uno.edu
Connection: Keep-Alive
```

save, then trim away
HTTP headers to
retrieve image

Use: e.g., WinHex

Save As Print Entire conversation (86631 bytes) ASCII EBCDIC Hex Dump C Arrays Raw

Filter Out This Stream Close



Conclusion: Network Analysis

Potentially a source of valuable evidence beyond that available from “dead” analysis


By the time an incident occurs, may have lost the chance to capture much of the interesting traffic

Challenging: huge volumes of data

Again, only one part of a complete investigative strategy

This introduction didn't include stepping stone analysis, many other factors (limited time)

Cloud Forensics



Cloud Forensics lies at the intersection among

- digital forensics
- computer network = access to resources is mediated by a network
- cloud computing

Forensics in Perspective : Incident Response



- A better perspective may be achieved if we consider incident handling
- Now forensics is considered in the Analysis step to discover which are the causes of the incident, who is involved etc.

B.Grobauer, T.Scherck, Towards Incidents Handling in the Cloud, CCSW, Oct. 2010



Issues, approach and challenges

For each of the four step in the following we discusses

- Issues = Problems to be faced
- Possible Approaches
- Challenges

In the following CSP = Cloud System Provider



Detection

- Timely detection of security incidents depends on systematic event monitoring.
- Event monitoring must be geared towards the detection of security incidents this requires that
 - all relevant existing event sources (e.g., OS and application logfiles) are monitored,
 - security-specific event sources (e.g., intrusion detection systems) are added where necessary,
 - adequate methods for identifying events that may indicate a security incident are utilized.



Customer Issues for Detection

No access to CSP-controlled event sources and vulnerability information

- The customer has no access to
 - events generated by infrastructure components that are under the control of the CSP
 - information about vulnerabilities found in the CSP-controlled infrastructure components.
- For PaaS, customers typically only have access to events generated by their own application (e.g., via application logging);
- For SaaS, customer completely depend upon the CSP for activity logging, etc.
- For IaaS
 - the problem is somewhat less acute, because virtual servers are under the customer control.
 - but the underlying virtualization infrastructure as well as parts of the network infrastructure connecting the virtual servers are controlled solely by the CSP.



Customer Issues for Detection

Insufficient interfaces for access to relevant data

- With for PaaS and SaaS,
 - access to information relevant for incident handling must necessarily occur via interfaces under the control of the CSP.
 - These interfaces may be insufficient for integration of the available data into event monitoring systems. For example, logging information displayed via a management web interface can be viewed by a user but is hard to process in an automated way.
- With IaaS, customers usually will be able to access event information from virtual servers in a way suitable for automated processing, but for all CSP-controlled data, the same problem as for SaaS and PaaS occurs..



Customer Issues for Detection

Inability to add security-specific event sources

- With infrastructure under one's own control (or with a customer-tailored offering), security-specific event sources can be added when required. For example, protect a web application through a web-application firewall.
- With cloud offerings, such additions become difficult if not impossible

Misdirection of abuses/incident reports

- Incident may be discovered thanks to third parties. In a cloud business model it is often unclear to whom abuse/incident reports should be directed.
- in IaaS scenarios, incident reports regarding abusive traffic from a certain IP address will be directed to the CSP rather than the customer whose virtual server has been causing the abusive traffic. It may be difficult for the CSP to find out, to which customer the report refers.
- In SaaS scenarios, reports regarding the compromise of a customer's SaaS application points to application weaknesses that affect any customer.



Detection Possible Approaches

To enable customers to reliably detect incidents, CSPs should adapt their service levels and offerings as follows:

- **Access to relevant data sources**

Considerations of which data sources are relevant for incident detection at the customer side must lead to service-level agreements that describe data sources an access possibilities.

- **Incident detection and reporting obligations / service**

- Incidents that originate with CSP-controlled infrastructure and might have an impact on a customer's resources must be reported to the customer.
- The SLA must provide a well-defined incident classification scheme and inform about reporting obligations and service levels (what is reported, how fast is reported, etc.)
- As an alternative or supplement to providing access to relevant data sources as described above, the CSP may offer an incident detection service that monitors thesedata sources for possible security incidents.



Detection Possible Approaches - 2

- **Open interfaces for event/incident data exchange**
Since systematic event monitoring is at the core of timely incident detection, the CSP should enable systematic event-monitoring by offering open/standardized interfaces for accessing event and/or incident data.
- **Intrusion-detection/prevention service portfolio**
 - Since customers usually cannot add intrusion detection/prevention capabilities to their cloud resources, the CSP may have to offer such capabilities, possibly as service add-on.
 - An alternative is to offer the integration of third-party service for intrusion detection/ prevention: it is to be expected that feasible service offerings in this direction evolve as “security-as-a-service” cloud offerings.
- **Acceptance and forwarding of external incident reports**
 - The CSP must accept external incident reports, ideally following established best practices and standards
 - External incident reports that concern or impact a customer must be brought to the attention of the customer with a defined service level.



Detection Challenges

- **Identification of relevant data sources**
 - It is not straightforward to determine, relevant data sources for incident detection.
 - In SaaS and PaaS, methodologies must be adapted to service paradigms: how can intrusion detection be carried out at the application level?
- **Standardization of event information**
 - No leading standard for expressing event information has emerged out of the field of existing initiatives
- **Customer-specific logging**
 - Because of “events generated by the infrastructure may concern
 - non-customer specific parts of the infrastructure,
 - resources of a single customer,
 - resources of several customers.
 - For providing customers access to event sources, the CSP must implement concepts and mechanisms that ensure that
 - all relevant event information should be accessible,
 - one customer should not be able to view event information regarding other customers.

The two goals may be conflicting for events concerning several customers



Detection Challenges

- **Detection in spite of missing information about customer infrastructure/resources**
- Security-services regarding intrusion detection and incident detection must take into account that the CSP has little or no knowledge about the customer infrastructure/resources.
- This problem is most pronounced with IaaS e.g. when providing intrusion detection for virtual machine images without knowledge regarding the installed OS
- It but also occurs with PaaS, e.g., intrusion detection for web applications without knowledge about the application.



Customer Analysis Issues

Limited knowledge about architecture

- The analysis of an incident, requires detailed information about network infrastructure, system configuration, and application-specific implementations
- For those parts of the infrastructure under control of the CSP, such information is usually not available because the exact set-up of the cloud infrastructure must be regarded as the CSP's core intellectual property.
- If a PaaS application accesses data from another service hosted at the same provider, for the customer, it is unclear how this access is implemented because the access occurs via an API call and no more information is provided.
- If an attacker subverts this mechanism and starts to redirect calls between applications, the customer may eventually notice that something is wrong and eventually detect a security incident.
- But the customer will find it very hard to analyse the incident. At best, the customer can suspect that something is wrong with intra-PaaS access to other applications but cannot to verify that assumption



Customer Analysis Issues

Missing knowledge of relevant data sources

- Ignorance about the architecture and infrastructure entails ignorance about data sources relevant for analyzing an incident.
- As an example, if the customer does not know about the PaaS-internal DNS service to resolve requests, he does not know about the log files of the DNS service that might be useful for understanding the security incident.

Unclear incident handling responsibilities

- As the two previous issues show, there is a clear need for co-operation between the CSIRT of the customer and some incident handling capability of the CSP
- For most current cloud offerings, there is no clear sense of the CSP's responsibilities in case of a security incident, let alone well-defined interfaces between the customer and the CSP in case of an incident.



Customer Analysis Issues

Problems of gathering evidence and forensics evidence

- In a traditional IT infrastructure evidence may gather by creating a 1:1 copy of the system's hard disc. With cloud computing, the situation is different:
 - Systems are out of the customer's reach
 - virtualized rather than physical
- For IaaS,
 - the virtual machine image can be attributed to exactly one customer, so handing over a 1:1 copy of that image to the customer is a possibility.
 - For event information, e.g., network firewall logs, the situation is more complicated, because they may include also other customer's information
- With PaaS and SaaS,
 - services are shared on one machine for several customers and it is not possible to give a 1:1 copy of the system to one customer.
 - the system and application logs often include information for several customers and cannot be easily accessed by the customer.



Analysis Possible Approaches

Provision of technical information about infrastructure

- When entering a cloud-sourcing relationship, cloud customers should have at least a basic understanding of the CSP's infrastructure such that in a security incident, information gathering does not “start from zero.” The CSP should provide such information to the customer.

Access to relevant data sources

- Considerations of which data sources might be relevant for incident analysis at the customer side must lead to appropriate SLA that describe access possibilities to such data in case of an incident.
- Alternatively, the CSP can analyze data according to the questions of the customer's CSIRT and provide the customer with analysis results.



Analysis Possible Approaches

Access to CSP incident handling capability

- The above-mentioned analysis of data sources is an example that customers may require access to the CSP's incident handling capability.
- The CSP's incident handling capability must have clear responsibilities regarding the co-operation in the analysis of security incidents that should be described in the SLA.



Analysis Possible Approaches

- **Interface to forensic use of virtualization technology**
- For IaaS, virtualization allows novel methods of carrying out forensic analysis which should be made available to IaaS customers.
 - It allows an investigator to introspect the compromised host. For example, Xen provides access to the runtime state of any virtual machine running via the Xen hypervisor (VM introspection).
 - virtual machines can create snapshots that can be taken and provided for forensic examination.
 - Snapshots are actually advantageous for forensic analysis, because an attacker cannot easily remove his traces on the system



Analysis Challenges

- **Separation of customer's data sources during evidence collection**
- As with data sources for detection, resource pooling causes data sources relevant for incident analysis to include data of many customers.
- When one customer is provided access to a data source, the CSP has to assure that this customer does not see information regarding other customers.
- **Adapting forensic analysis methods to the cloud**
- Current forensic methods are geared towards traditional IT infrastructures. It is unclear how to effectively perform incident analysis in a highly dynamic cloud computing environment with redundancies data mobility.
- Attacks are changing for the cloud e.g. botnets will use cloud computing to hide their activities and new methods are developed to detect and analyze such kind of attacks.
- First steps towards improving incident analysis should improve
 - live analysis techniques
 - improving log file analysis.



Analysis Challenges

- **Improving live analysis techniques**
- Live analysis such as memory forensics is fundamental and it is an example of how a snapshot of a virtual machine image could be analyzed.
- Forensics must be performed on the running system anytime a snapshot is not available. Such an approach
 - carries certain risks – if the attacker has completely subverted the system, he has the means to hide his activities very effectively
 - in many cases valuable information can be learned.
 - Proprietary approaches towards streamlining live analysis but a comprehensive approach is currently lacking.
- **Improving log generation & analysis techniques**
- The importance of log file analysis rises in cloud computing especially for PaaS and SaaS, because most available information about an incident will be contained in log files.
- It is therefore essential to improve
 - the generation of logging information
 - analysis techniques for logging information



Containment, Eradication, and Recovery

IaaS

- A frequent incident scenario in an IaaS setting is that a virtual machine image has been compromised by an attacker.
- A common first containment step in the corresponding non-cloud setting – the compromise of a server – limits or cuts network connectivity, where “limits ranged from blocking communication with certain network parts to routing traffic via an active device in order to observe and selectively block traffic.
- Which of these activities can be carried out in a cloud-setting heavily depends on the network configuration capabilities offered by the cloud provider.



Containment, Eradication, and Recovery

IaaS

- For a different course of action, the cloud setting actually is very beneficial: virtualization offers the possibility to “pause” a vm image, which at the same time blocks further attacker activities and preserves full information for further analysis.
- The elasticity feature of provisioning a vm with more or less resources according to demand may be useful in containing an attack →
 - a compromised vm can be easily starved of resources, thus slowing down attacker activities such as abuse of the compromised system to send spam or attack other systems
 - in some scenarios, adding resources via elasticity may be helpful in mitigating a DoS attack.



Containment, Eradication, and Recovery

IaaS

- Also eradication and recovery may be aided by the cloud setting: if the point of time when the compromise occurred can be established, the snapshot feature could be used to revert the compromised virtual machine image to a non-compromised state.
- Such an approach, depends on well-established change management processes such that legitimate changes to the virtual machine image after a snapshot has been taken are tracked.



Containment, Eradication, and Recovery

IaaS

- In the previous scenario, CSPs can help their customers in containment, eradication and recovery by offering the following:
 - **Ability to configure networking** The more flexible the network configuration, the more options for containing an incident exist.
 - **Access to halting and snapshot features of virtualization** By providing a “snapshot and restore facility” to the customer, eradication and recovery activities can be supported.
- In scenarios where the attacker exploit a vulnerability in the underlying infrastructure of the CSP, the customer cannot contain, eradicate and recovery for the virtual machine images hosted by that CSP
- Once the IaaS market has matured some more so as to allow easy transfer of virtual machine images between providers, moving virtual machine images from a compromised provider may become a possible option to start the process of containment, eradication and recovery.



Containment, Eradication, and Recovery

SaaS and PaaS

- Most attacks are enabled by application vulnerabilities and the attacker can compromise accounts and/or elevate privileges of some accounts
- Containment essentially reduce or completely remove functionality that allows the attacker to carry out unauthorized activities
- If these functionality cannot be restricted, an alternative may be to closely monitor the functionality and then timely react to abuse.
- Depending on the scope of the vulnerability and the capabilities for reducing functionality the whole application has to be taken offline to make some adjustments in functionality.
- A work around when dealing with web applications is to use web application firewalls to close known attack vectors until the root cause can be treated.
- The ability for containment in a PaaS and SaaS setting depends upon
 - the granularity with which functionality and access rights can be configure
 - the ability to implement work around



Containment, Eradication, and Recovery

IaaS, SaaS and PaaS

- For eradication and recovery, the application vulnerability has to be closed.
 - for SaaS, this is clearly the obligation of the CSP,
 - for PaaS it depends whether the vulnerability lies in the customer's code or the implementation of API functionality provided by the CSP and used in the customer's code.
- To eradicate & recovery the customer must purge its data in the application from the attacker's activity that may have uploaded malware-infected content or modified existing content. This requires
 - precise logging information of all data changes
 - direct administrative data access



Containment, Eradication, and Recovery

IaaS, SaaS and PaaS

In any incident scenario CSPs can help their customers in containment, eradication and recovery by offering

- **Granular configuration of functionality and access rights**

The more granular the configuration of functionality and access rights is, the higher the chance that vulnerable features in a SaaS application can be disabled in a limited scope that contains the incident but allows continued use of the application.

- **Possibility to configure web application firewall**

If the CSP offers the customer the possibility to configure a web application firewall for his PaaS applications, it may be possible to carry out containment using detection and prevention possibilities of the web application firewall.

- **Direct read/write access to customer data**

A direct read/write access to customer data rather than only via the application GUI, simplifies eradication and recovery at the data level for the customer.



Implications

Identify relevant event sources

- The customer needs to identify possible approaches to detect and analyze security incidents.
- The most important basis for analysis and detection is event information – therefore, relevant sources of the cloud service under consideration and the possibilities to add security-specific event sources must be identified.

Evaluate CSP's level of support for detection and analysis

- Does the CSP provide access to the relevant event sources?
- Are the CSP's own incident handling capabilities adequate?
- Are incidents that have been detected by or reported to the CSP communicated in a timely fashion?
- Does the CSP provide adequate access to information the analysis requires?

Establish communication channels and exchange formats

- The customer relies on the access to event information and incident reports as well as efficient communication with the CSP's incident handling capability for analysis and response.



Implications

Formats used for communicating event and incident information:

- The Incident Handling tools at the customer's side such as incident tracking system and tools for event analysis must be able to work with the formats used by the CSP.

Evaluate interface to contain, eradicate, and recover an incident

- Probable incident scenarios suggest that standard mechanisms, such as
 - customer access to virtualization snapshot functionality,
 - customer-configurable web application firewalls, customer
 - direct data access

can be helpful for containment, eradication, and recovery.

- The wide range of possible incident scenarios implies that standard mechanisms will be not sufficient in all cases.
- Therefore, in many cases adequate access to incident handling personnel of the CSP will be essential.



Open Issues

Cloud SLAs for incident handling.

- Several issues regarding incident handling should be treated in cloud SLAs.
- Precise requirements on CSPs regarding incident handling must be defined and included in standard cloud security requirements such as the
 - Common Assurance Maturity Model
 - Cloud Security Alliance's trusted cloud
- For any outsourcing there is a SLA.

Generating and processing event information

- Every cloud environment is different and requires a systematic approach towards to identify relevant events to detect and analyse attacks
- Efficient handling of event information requires accepted standards for event information.
- Resource pooling leads to event sources with information about many customers that cannot be made accessible to a single one for incident detection and analysis: methods for generating customer specific logs that do not violate the confidentiality/privacy requirements of other customers are required.



Open Issues

Detection and Analysis

- Virtual-machine introspection is uniquely suitable for incident analysis in a cloud; further research about virtual-machine introspection and its use for incident handling in the cloud should be conducted.
- To make the most of virtual-machine introspection and snapshots, research in memory forensics must be intensified.
- The collection of information via live forensics on running systems must be subjected to a systematic approach.
- Methods for detection and analysis based on event information such as logfile correlation and visualization must be improved and adapted to incident handling in the cloud. This is of special importance for incident handling in PaaS and SaaS, where most relevant information will be available as event logs.
- Detection methods that require little or no information about the monitored infrastructure (e.g., virtual machines under customer control or web applications under customer control, anomaly-detection approaches to web-application firewalling) must be improved.



Another point of view

Cloud forensics: An overview

Keyun Ruan, Prof. Joe Carthy, Prof. Tahar Kechadi, Mark Crosbie

This paper classifies

- Cloud challenges problems w.r.t. traditional solutions
- Cloud opportunities advantages w.r.t. traditional solutions



Challenges

forensic data collection

- In all combinations of cloud service and deployment models, the cloud customer has decreased access to forensic data.
- IaaS customers enjoy relatively easy access to all data required for a forensic investigation, while SaaS customers may have little to no access to data required
- Decreased access to forensic data means the cloud customers generally has no control or knowledge over the physical location of their data, and may only be able to specify location at a higher level of abstraction, typically as an object or container
- CSPs intentionally hide the location of data from customers to facilitate data movement and replication.
- Moreover, there is a lack of appropriate terms of use in the SLA to enable general forensic readiness in the Cloud.
- Several CSPs do not provide services or interfaces for the customers to gather forensic data.



Challenges

elastic, static and live forensics

- The proliferation of endpoints is a challenge for data discovery and evidence collection simply because of the number of Cloud resources
- Accurate time synchronization is both crucial and challenging as physical machines spread in multiple geographical regions must be synchronized
- Log formats is a traditional issue in forensics and it is exacerbated in the Cloud because it is extremely difficult to unify formats or make them convertible to each other
- Recovering deleted data is an important source of evidence even in the Cloud.
- In AWS the right to alter or delete a snapshot is explicitly reserved for the account that created the volume. When item and attribute data are deleted within a domain, removal of the mapping starts immediately, and is also generally complete within seconds.
- Storage space occupied by delete elements is made available for future write operations and it is likely that it will be overwritten by new data.



Challenges

evidence segregation

- On the physical level system audit logs of shared resources and other forensic data are shared among multiple tenants.
- Currently, it is a challenge for the CSP and law enforcement to keep the same segregating in the whole process of investigation without breaching the confidentiality of other tenants sharing the same infrastructure and ensure the admissibility of the evidence.
- Easy-to-use feature of cloud results in a weak registration system, facilitating anonymity that is easy to be abused and making it easier for cloud criminals to conceal their identities and harder for investigators to identify and trace suspects as well as segregate evidence.



Challenges

evidence segregation

- Encryption is used to separate data hosting of the CSPs and data usage of the cloud customers and most CSPs encourage customers to encrypt their sensitive data before uploading to the Cloud as unencrypted data in the Cloud can be considered lost from a strict security perspective.
- A chain of separation is required to segregate key management from the CSP hosting the data and needs to be standardized in contract language.
- Agreement has to be made among the law enforcement, the cloud customer and the CSP on granting access to keys of forensic data, otherwise evidence can be easily compromised when encryption key is destroyed.



Challenges

virtualized environments

- Most cloud environments exploit a virtualized environment, monitored and provisioned by a VMM
 - Attackers will aim to focus their attacks against the hypervisor
 - lack of policies, procedures and techniques to facilitate investigation on VMMs.
 - data it is stored in multiple jurisdictions and the lack of real-time information about location introduces difficulties for investigation.
- Investigators may unknowingly violate regulations, especially if clear information is not provided about the jurisdiction of storage
- The CSPs cannot provide tools for the customer to locate at a given time, or trace at a given period of time, precisely and physically the multiple locations of a piece of data across all the geographical regions
- The agency of a single nation cannot manage cases such as confiscating “a Cloud” if the physical servers are spread across different countries.



Challenges

internal staffing

- Most cloud organizations are dealing with investigations with traditional network forensic tools and staffing
- The major challenge in establishing a cloud forensic structure is the lack of forensic expertise and relevant legal experience.
- The deep-rooted reasons for this challenge are
 - the relative slow progress of forensic research compare to the rapidly evolving technology
 - the slow progress of relevant laws and international regulation
- Digital forensics is still in its infancy, new research areas in non-standard systems such as cloud computing, need to be explored, techniques need to be developed, regulations need to catch up, law advisors need to be trained, staff need to be equipped with knowledge and skills grounds



Challenges

external chain of dependency

- CSPs and most cloud applications often have dependencies on other CSPs. For example, a CSP providing an email application (SaaS) may depend on a 3rd party provider to host log-files (PaaS), who in turn may rely on a partner to provide infrastructure to store log files (IaaS).
- Although many predict the industry is moving towards federated or integrated Cloud in the near future, today every CSP has a different approach to solving this problem. Correlation of activities across CSPs is a big challenge.
- Investigation in the chain of dependencies between CSPs may depend on the investigations of each chain link and level of complexity of the dependencies.
- Any interruption or corruption in the chain or a lack of coordination of responsibilities between all the parties involved can lead to problems. Currently there are no tool, procedure, policy or agreement regarding crossprovider forensic investigations.



Challenges

SLA

- Important terms regarding forensic investigations are not included in the SLA due to a lack of
 - customer awareness,
 - CSP transparency
 - international regulations.
- Most cloud customers are still not aware of the potential issues regarding forensic investigations in the Cloud and their significance. Hence, they might end up not knowing anything at all about what has happened in the Cloud in cases when their data is lost in criminal activities and has no right to claim any compensation.
- CSPs are not willing to ensure transparency to the customers regarding forensic investigations because they either do not know how to investigate cloud crimes themselves or the methods and techniques they are using are likely to be problematic in the highly complex and dynamic multi-jurisdiction and multi-tenancy cloud environment.
- The progress of any law and regulations including law and regulations of cyber crimes is very slow, while cloud computing is rapidly emerging as a new battlefield of cyber crimes for hackers who are equipped by the most updated techniques, investigators, law enforcement and various cloud organizations.



Challenges

Multi-Jurisdiction and multi-tenancy

- The legal challenges of multi-jurisdiction and multi-tenancy concern the differences among legislations in all the states the Cloud and its customers reside in.
- The differences between jurisdictions affects on issues such as
 - what kind of data can be accessed and retrieved in the jurisdiction(s) where the physical machine(s)
 - which data is accessed and retrieved,
 - how to conduct evidence retrieval without breaching privacy or privilege rights of tenants according to the privacy policies and regulations in the organizations
 - specific jurisdiction where multiple tenants' data is located,
 - what kind of evidence is admissible to the court in the specific jurisdiction,
 - what kind of chain of custody is needed in the evidence preservation in the jurisdiction(s) where forensic data has passed during an investigation in the Cloud.
- Multi-jurisdiction issues also concern lack of legislative mechanism that facilitates collaboration between industry and law enforcement around the world, in cases such as resource seizure, cloud confiscation, evidence retrieval, data exchange between countries, etc.



Opportunities

Cost Effectiveness

- Everything is less expensive when implemented on a larger scale, including security and forensics
- The cost advantages of cloud computing applies to forensics. SMEs that cannot afford dedicated internal or external forensics implementations or services may have an upgrade at relatively low cost when adopting cloud computing.

Data Abundance

- Clouds ensure object durability by multiple copies across multiple Availability Zones on the initial write and then actively doing further replication in the event of device unavailability or detected bit-rot to reduce the risk of single point of failure+
- Data abundance is helpful to investigations as full data deletion cannot be guaranteed and investigators can take advantage of it to recover data as evidence.
- When a request to delete a cloud resource is made it actually technically can never result in true wiping of the data. Full data deletion may only be guaranteed by destroying the resource shared with other cloud tenants. Thus pieces or segments of data crucial to investigation are very likely to remain somewhere in the Cloud for the investigators to discover.



Opportunities

Overall Robustness

- Cloud technologies help to improve the overall robustness of forensics IaaS offerings support on-demand cloning of VM. In the event of a suspected security breach, the customer can take an image of a live VM for offline forensic analysis, leading to less downtime for analysis
- Multiple clones can be created and analyzed in parallel to improve the analysis of incidents and increase the probability of tracking attackers and patching weaknesses.

Scalability and Flexibility

- Cloud computing allows scalable and flexible usage of resources which also applies to forensic services.
 - It can provide unlimited pay-per-use storage of logs, allowing more comprehensive logging without compromising performance. It can also increase the efficiency of indexing, searching and various queries of the logs. Cloud instances can be scaled as needed based on the logging load.
 - Forensic activities only take place when incidents happen which can largely take advantage of the cost-effectiveness of cloud computing.
 - Customers have the choice to build their own dedicated forensic server(s) in the Cloud, ready to use only in need.
-