

After an Attack



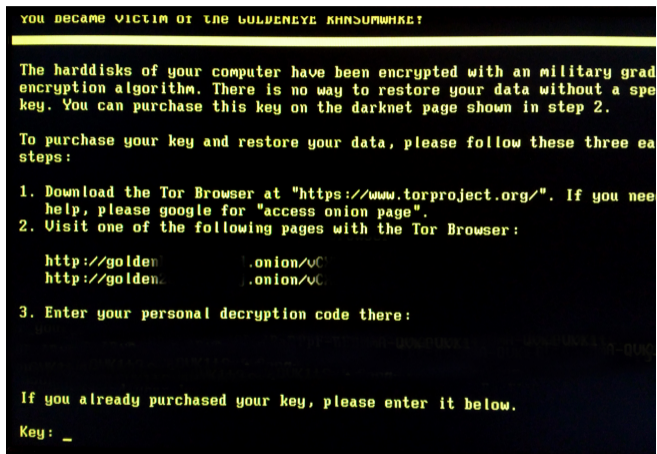
After an Attack

- Detection
- Non-technical recovery
- Technical recovery

Detecting Successful Attacks

- How do you know if you've been hacked?
- Many attacks are not noticed for a long time
 - Equifax: Four months
 - OPM: 11 months
 - Yahoo: Two years
 - Marriott/Starwood: Four years
- What took them so long?

Some Attacks Announce Themselves...



(Image by jbuket; taken from <https://en.wikipedia.org/wiki/File:PetyaA.jpg>)

This is probably bad news

Announcement May be the Point!

- The Sony hack: North Korea wanted to dissuade them from releasing a particular movie
- The DNC and Podesta hacks
- Hacktivism in general

Third Party Detection

- Often, what can easily be seen is the consequences of a hack
- Who the third party is varies

Credit Card Numbers

- Credit card issuers track stolen cards and look for common prior uses
- If they find a cluster, they investigate—were they hacked, or is that merchant itself evil?
- They then do the appropriate notification
- This is how Target found out it had been hacked

(Detecting Stolen Cards)

- First, of course: customer reports
- Second: they (of course) use machine learning these days
- Look for suspicious patterns: location of use, high dollar value (especially for electronics or jewelry), often preceded by a low-value “probe” transaction
- Other things they won't tell us. . .

Watching Underground Sites

- Often, stolen information is sold via underground markets
- Watch those sites to see if some of your data is for sale
- That's how Yahoo learned it had been hacked two years earlier: its credentials were for sale
- This can also be used for confirmation: banks bought groups of stolen credit card numbers and found that many had been used at Target and (later on) Home Depot

- Marriott's "security tool alerted Marriott officials to an unauthorized attempt to access Starwood's guest reservation database."
- Why doesn't that always work?
- (Why did it take Marriott 4+ years? Did Marriott's IT group have better tools than Starwood's pre-acquisition group?)

Underground Sites as an Attribution Clue

- The data stolen from Marriott did not show up on underground markets
- “Usually when stolen data doesn’t appear, it’s a state actor collecting it for intelligence purposes”
- From the first lecture of the semester: *What are you trying to protect against whom?*

Internal Investigations

- The IT staff at OPM (Office of Personnel Management in the US government) was investigating encrypted traffic
- There was some going from what appeared to be a McAfee module to `opmsecurity.org`—but OPM didn't use McAfee software and `opmsecurity.org` wasn't used by them
- In fact, it had been registered a year earlier
- Oops...

- Most sites do not notice promptly that they've been hacked
- Why not? Mostly, because they don't look
- A better assumption: your site *has* been hacked and is penetrated *right now*
- The only solution is internal intrusion detection: detectors that find suspicious activity inside your network
- Many types—take COMS E6185...

Exfiltration Detection

- If your primary asset is bulk data, look for unusual outbound flows
- Sony reportedly lost 100 *terabytes* of data
- Equifax (by my estimate) lost > 30 terabytes
- Why didn't anyone notice the anomalous traffic flows?
- (But: what about “low and slow” exfiltration”?)

Non-Technical Recovery

Many Aspects

- Forensic examination
- Criminal prosecution
- Public relations
- Personnel
- Product operations
- Legal
- Often more

Have a plan and rehearse it ahead of time

Who Runs Recovery?

- Legal? Perhaps important in regulated industries
 - 👉 Leak of personal information is almost always a legal issue
- IT? Possibly, in a small shop; larger ones it might be the CSO
- Who controls what PR says?
- What do you tell employees? How do you prevent things from leaking, if you don't want to alert the attackers you've found them?
- Again: plan ahead of time

- Shutting down a network to disinfect it interferes with work, which interferes with operations
- How much interference is acceptable?
- How can it be scheduled?
- But: attackers can also disrupt operations
- The balance is probably a CEO- or board-level decision
- In a serious hack, e.g., ransomware or data destruction, how do you operate during recovery?

Operations During Recovery

- Operate offline if you can (whiteboards! employee cellphones!— VoIP phones and smart doorlocks can also be affected)
- Who handles payroll?
- Order new computers? It's been done!
- Email?

- What do you disclose?
- (Some disclosures are legally mandated, including (sometimes) to investors)
- Does disclosing more or less information help or hurt the company's reputation?
- (Might you be hauled before a Congressional investigating committee? It happened to Equifax management. Correction: to their *former* management; senior executives, including the CEO, lost their jobs because of the breach.)

- Suppose you want to prosecute the bad guys
- Should you do these forensics?
- No! Forensic analysis has become highly specialized (take COMS W4186)
- But—you *must* know enough about what happened to be sure you can expel the intruders and keep them out

- Evidence must be handled *very* carefully
- Must avoid defense charges of tampering, forgery, misinterpretation (to say nothing of legal issues such as proper warrants)
- Parties with more interest in a case can be portrayed as biased

- Chain of custody
- Disk copies made using specialized hardware to prevent accidental overwrites
- Rigorous marking, labeling, logging, etc.
- Careful records of all analysis
- Not a job for amateurs

Calling in the Police

- Establish relationships with law enforcement before you have a problem
- Find the right agency or the right officers—many police forces, especially smaller ones, don't have the right expertise
- Learn whom to call
- Learn what *they* want you to do
- Learn about specialized threat information sources for your industry

International Dimensions

- Many—most?—attacks cross national boundaries
- Warrants good in one country are not valid abroad
- The available legal process, MLAT (Mutual Legal Assistance Treaty), is slow and cumbersome
- Some national police forces have good relationships with others, especially in the EU and among the “Five Eyes” countries

Technical Recovery

Forensic Assessment—Why?

- What has to be thrown out?
- What can be saved?
- How did the bad guys get in?
- How do you get them out of your network and keep them out?

A General Rule

- It is frequently impossible to cleanse an infected system
- Hiding back doors is relatively easy
- The usual advice: reformat your disks and reinstall
- Also: firmware in peripherals—disks, keyboards, USB devices, and more—can be replaced with malware. Maybe you need to discard the hardware! (This is sometimes done!)

Recovery Goals

- Get back on the air
- Prevent reinfection
- Find *all* infected machines
- Do this with high assurance

Assurance!

- It's not just fixing things, it's *knowing* that you've fixed them
- After all, have you fixed the problem—or do you merely think you have
- How will you behave if you think your system might be compromised?
- How will you behave if it really is secure?
- 👉 What if you behave as if it's secure, but it isn't?
- Assurance is *knowing* the actual security status

Tolerating an Attacker

- Sometimes, the best way to find the infected machines is to tolerate the attacker
- See how the infected machines communicate, and how they go after other machines
- But: you can't get caught watching them
- And: you risk further damage
- Are there liability issues?

Attacker Patterns

- Initial entry
- Scouting
- Privilege escalation
- Lateral movement

- Look for any of those
- Correlate with employee behavior, e.g., was an employee logging in while on vacation?
- Off-machine detectors are especially valuable, to avoid alerting the attackers

- Check your logs for suspicious entries from compromised machines

 Netflow data is especially important here

- Outbound connections from known-infected machines can indicate attempts to spread the problem
- Earlier, inbound connections *to* the infected machine can show how the problem started—and identify other infected machines
- (If the infection came from outside, do you notify the site? How?)

Backups Are Your Friend

- Back up your system frequently
- Make sure you have a 0-day backup, from before the system went live
- Recover your data—but not your programs—from the backups
- (Also: *test* your backups frequently, and test your recovery processes. Also, keep some backups offsite, as protection against fires, etc.)

Backups Are Your Enemy

- Your system was infected, or at least vulnerable, “yesterday”
- If you just do a restore, you restore the problem
- Partition your backups

Partitioned Backups

- Generally, don't back up the operating system or external applications—on restore, use the newest version from the vendor
- But: do back up local configurations and locally written software
- Back up databases separately—most (but not all!) won't carry infections (but watch out for credential and authorization databases!)
- Emails...

Where the Bad Things Are

- Persistent infections are likely to linger in configuration files, credential and authorization databases, software packages—and email
- How do you issue new credentials to tens of thousands of employees, throughout the world?
- How do you audit authorization and configuration databases?
- What about infected email attachments?
- Note well: dealing with any of these things requires planning ahead

Centralized versus Decentralized System Architectures

- In a centralized system, it's easy for a central group to push out fixes everywhere
- It also makes it easy for an attacker to spread everywhere and to do damage everywhere
- Example: Maersk shipping line and NotPetya—there was one (replicated) Active Directory domain controller for the entire company, which NotPetya destroyed
- The only surviving copy was on a temporarily offline machine in Ghana—they had to hand-carry it to Nigeria and thence to London (visa issues!)


Decentralized Architectures

- Different groups will do things differently—no central point of failure, but no central point of repair
- Does every group do things securely?
- In the Equifax hack, 1 of 117 web servers wasn't patched
- Lateral movement is still possible—different parts of the company talk to each other, some employees cross domains, etc.
- Recovery is harder—you can't clean up everything at once
- Recovery is easier—you can clean up sections at a time, and try to isolate them
- N.B.: system architectures tend to follow overall organizational structure of the company

Ransomware

- Malware that encrypts your files—you have to pay to get the decryption key
- Some of your files may be leaked, both for embarrassment and as proof that the attacker was in your system
- If you pay, you get the decryption key; if not, you better hope you have good backups
- Restoring everything can take a *long* time—can you afford to be off the air that long?
- Some types of ransomware encrypt your backup files, too—keep a copy offline, if you can (but that's hard for large sites)
- (Some ransomware isn't decryptable; it's just vandalism plus payment)

- Again: this is a job for experts; all I'll do is give an overview
- *Never* try to work with a live disk
- You don't want to destroy metadata
- Be careful of the malware!
- Make a copy an image copy
- Don't use anything that will change file access times
- Free space can be important

 This is what law enforcement does when analyzing seized computers

- If you don't have a spare machine (with compatible hardware), trying booting a "live" CD or USB stick
- A live disk is a bootable, runnable system
- Example: Ubuntu installer; TAILS; MacOS installer

Mounting the Image

- Always mount it read-only, with the “noexec” and “nodev” options
- Most newer systems allow you to mount a file as a block device (vnd on BSD; lofiadm on Solaris; loopback device on Linux; .dmg files on Mac OS; etc.)

Things to Look For

- What files were changed recently?
- Note: look at `ctime`, not just `mtime` (why?)
(Windows has a similar set of file times.)

Finding Deleted Files

- Deleting a file doesn't delete the data
- Instead, it changes some metadata—the filename on FAT and NTFS filesystems; the i-node number and i-list entry on traditional BSD filesystems
- The blocks are returned to the freelist—but they may not be reallocated immediately
- Clever tools can recover deleted files

Rebuilding Deleted Files

- Suppose there are no clues in directories or the i-list
- Sometimes, it's possible to do magic with the freelist
- Files aren't random. . .

- Different file types have different byte distributions
- Example: C has lots of { and }; text has distinctive capitalization patterns, etc.
- Sort blocks by (probable) type

Contact Probabilities

- Look for matches between the end of one block and the start of the next
- Look for syntactically correct statements
- Log files have timestamps!

Are Deleted Files *Better* for Forensics?

- A normal file can be overwritten easily
- A deleted file can't be touched
- Block allocation policies are invisible to the application
- Some claim that deleted files are *more* likely to be intact

Looking at Memory

- If the system is still up, dump main memory (/dev/kmem)
- Can often find plaintext of the malware
- Encrypting file systems write ciphertext to disk—but where's the plaintext?
Often, in RAM
- The decryption key is in RAM, too, if the file system is mounted

Attribution

- This is also a job for experts
- Attribution is difficult, but not impossible (and it's easier today than it used to be—don't believe people who say it can't be done)
- You *cannot* rely just on network connectivity—attackers always use stepping stones
- Attackers are usually identified by a code name, e.g., Fancy Bear, Electric Panda, Equation Group
- Sometimes, there will be an (asserted) association with some country, especially if it's believed to be an intelligence group

- Look for common tools, tactics, command-and-control servers, Bitcoin wallets
- “False flag” operations are possible but harder to manage than some would have you think
- Note: having a library of previous incidents—their tools, tactics, C&C servers, etc.—is vital for technical attribution, which is another reason it’s a job for pros

Attribution: Non-Technical

- Use “all-sources intelligence” (if you’re in an intelligence agency)
- Use “open source intelligence” if you’re not
- Open source intelligence: strategically collecting and analyzing public but obscure information
- Examples: domain name registrations, email addresses, analyses of earlier hacks, etc.
- (See the Mandiant report on APT1 for detailed examples)
- Again: having a library of previous incidents is vital

Why Do Attribution?

- Learning who the attacker is tells you something about goals and abilities
- Example: Chinese and US intelligence don't try to steal money; North Korea does
- Example: intelligence agencies might rely on RAM-resident malware (i.e., it won't be on disk) and might infect device firmware

Recovering from an Attack

- There's a lot to do, especially in a large organization
- You have to do it, though, or your digital assets will remain at risk
- (Nortel tried for *10 years* and couldn't get Chinese intelligence out)
- Planning, rehearsals, tool-building, and preparations are key

Questions?



(Golden-crowned kinglet, Central Park, April 2, 2022)