

Security "Monsters:" Current Security Threats and What You Should Be Doing to Address Them

IT Security: A Call to Action for the Education Community
10:45-12:00, November 2nd, 2006, Ramada Plaza, Fargo ND

Joe St Sauver, Ph.D.
235 Computing Center
University of Oregon, Eugene OR 97403
joe@uoregon.edu or 541-346-1720

<http://www.uoregon.edu/~joe/monsters/>

0. Introduction

A Note About the Format of This Talk and a Disclaimer

- I've prepared this talk in some detail so that:
 - it can be followed by those who may not be present when the talk was originally given,
 - to insure that the contents of the talk are available to those in the audience who may be hearing impaired, and
 - to minimize the need for audience members to jot down notes.

Having a talk that's prepared in some detail also helps keep me on track.

- **Disclaimer: all opinions expressed in this document are strictly my own. Independently assess and reconfirm all recommendations presented, and note that even if you follow all recommendations given here, you may still experience a security breach.**

'Tis the Season For Ghosts and Goblins, Demons and Monsters!

- Unfortunately, some IT security "monsters" are real and will be there after the candy's all been handed out and the jack o' lanterns are gone.
- What we're going to do today is talk a little about some of those "monsters," and what you can and should be doing about them.

Today's "Monster Lineup"

- The Most Publicized Monster:
Compromised Personally Identifiable Information (PII)
- The Most Frequent Monster:
Malware (Viruses, Worms, Trojan Horses, Spyware, Root Kits, etc.)
- The Monster That Can Bite The Hardest:
Distributed Denial of Service Attacks, Spoofed Traffic,
BCP 38 Filtering and Open Recursive Name Server Abuse
- The Sneakiest Monster:
Address Space Hijacking

1. The Most Publicized Monster: Compromised Personally Identifiable Information (PII)

"The majority of higher education managers experienced at least one information technology security incident last year and one-third reported a data loss or theft."

Most campuses report security breaches, Oct 10, 2006

<http://www.fcw.com/article96412-10-10-06-Web&printLayout>

If a PII Spill Can Cause Someone To Get Fired, It's a Monster, Right?

- The current security "monster" that most keeps CIO's and Security Officers up at night is probably **unauthorized access to personally identifiable information (PII)**.
- Unauthorized access to PII data *ISN'T* actually the biggest IT security threat you face, however it *IS* **PERCEIVED** to be one of the most important issues you may be facing. (For a nice reality check, see "The Identity Theft Scare," *Washington Post*, Saturday October 14th, 2006, page A21.)
- Moreover, PII-related breaches have resulted in people getting fired (e.g., see: <http://thepost.baker.ohiou.edu/articles/2006/10/05/news/15373.html>)
- In any event, since PII is on people's minds, let's talk about it a little. What do we mean by "unauthorized access to PII"?

Some Examples...

- An administrative system containing credit card numbers has a privileged (“root”) login from Eastern Europe
- A faculty member gets a dataset from an insurance company with detailed patient records; sometime later that researcher's PC gets hacked & used as a warez site
- A laptop with student financial information is lost (or stolen from) a financial aid counselor who's travelling
- A “backup” CD with sensitive information can't be found.
- Clear text wireless network traffic ends up getting sniffed
- Grades & SSNs get posted on a (not so) "private" webpage
- A desktop is sold as surplus without its storage media being pulled and sanitized first; sensitive data gets extracted.
- An insider accesses private data, and sells that information to unauthorized people.

Nothing Really New Here...

- System compromises? They've been with us forever.
- Lost or stolen laptops or CDs? Not a new problem.
- Eavesdropping on network traffic? A longstanding risk...
- Remember when faculty members would post lists of last-four-digits of student IDs and final grades (in alphabetical order) on their doors? Pity poor Mr. Zzyniski... no privacy.
- Careless property transfer or negligent surplus property disposal practices? Veritably the stuff of legends.
- Untrustworthy insiders: yep, there have been those...
- And of course, loss of PII can even occur in non-IT formats (e.g., printed credit reports or charge slips are thrown in a dumpster unshredded)

So What IS Different About PII- Related Incidents Today?

- PII serves as an "**aggravating**" factor multiplying the gravity of formerly routine incidents
- Potential compromises receive "worst case" handling and end up being treated as if they're verified events (a cracker may not even know what he/she has accidentally "stepped into" and may not care about PII on a cracked system)
- Large numbers of people may be impacted by PII events (today's datasets are large, and may routinely contain hundreds of thousands or millions of records)
- More cracker/hackers are now monetarily motivated.
- Today people understand that PII breaches have non-remediable impacts (you can't "unring the bell"), and people have become acutely sensitized to the problem of PII disclosure in general.

Folks Feel Powerless When It Comes to PII Data Spills

To understand user sensitization, understand user psychology:

- PII spills have “unbounded” potential abuses (imagination may run wild and conceive "worst case" scenarios), often resulting in potential over-reaction relative to actual exposure.
- PII breaches often involve shadowy/unknown attackers who may be working from overseas, completely out of reach and with unknown motives ==> feelings of powerlessness.
- Discovery of a breach may be delayed ("What? The system with my data has been hacked for *eight months* and we're just finding out NOW?") ==> feeling of powerlessness.
- Unbounded exploitation time frames (just because they didn't steal your identity so far doesn't mean that they won't get around to it eventually) ==> feelings of powerlessness.

PII and Powerlessness (continued)

- Users feel an inability to take personal action to avoid exposure (dang hard to live & do business in the USA today w/o having a driver's license, credit cards, health insurance, internet access, etc.) ==> feelings of powerlessness.
- At the same time, you have no way individually of assessing the quality of the data stewardship at any given company; you just need to "trust them." Even if you trust one company, your data may be exchanged with other less trustworthy entities without your knowledge and explicit consent
- PII involves some key high anxiety areas for some people: money (in the case of financial data), their bodies (in the case of health related data), and/or their reputation/identity (in the case of other private data)
- At least some ID theft incidents have been highly publicized. Some may ask, "Why are PII events so newsworthy?"

Breaches of PII Are "Highly News-worthy Events" Because...

- **Events which inspire feelings of powerlessness are inherently newsworthy.** Consumers always seek info about those incidents to try to reduce feelings of powerlessness.
- Large number of real or potential victims ==> newsworthy
- Local example of a major national problem ==> newsworthy
- Media-perceived duty to educate potential victims about their exposure ==> newsworthy
- "Inherently scandalous event" inspiring moral outrage at apparent negligence and/or delivering an opportunity to promote reform? Newsworthy!
- Opportunity to lampoon apparent institutional incompetence ==> **VERY newsworthy**
- Bottom line? **PII incidents WILL get covered by the media**

News Coverage Drives Other PII-Related Phenomena

- News coverage may make some institutions want to minimize disclosures regarding an incident, which gives the impression that there's a "big secret" to be ferreted out (which causes newsworthiness to increase further... a vicious, ugly, cycle).
- News coverage leads to public pressure to "do something" about the problem; that public pressure gets translated into political attention which leads to new laws concerning PII breaches. Future incidents become a bigger deal still.
- News coverage can potentially interfere with ongoing law enforcement (LE) investigations ==> the bad guy/gal doesn't get caught AND the attacks/problem continues AND no intelligence about the originally stolen PII gets obtained.
- News coverage may lead to scapegoating and victimization of innocents who weren't actually responsible for the breach.

Practical Steps

Protecting PII, when you get right down to it, is about good general IT security practices; the presence of PII is just an aggravating factor layered on top of what would otherwise be a "normal" IT security incident.

Prudently planning for worst case scenarios and recognizing the value of a layered defense, what can be done to avoid PII disclosure or at least minimize the aggravating impact of PII if a breach were to occur?

Attend to Basic Desktop & Server "Bread and Butter" IT Security Issues

- Antivirus and antispyware protection running and current
- OS (and all applications!) are patched up to date
- Software firewall running
- Systems routinely backed up
- Less vulnerable applications selected and deployed
- System hardened (e.g., unneeded system services disabled; no files shared/exported; etc., etc., etc.)
- Routine day-to-day use of non-administrator accounts
- Strong, periodically changed, passwords
- All data on disk and all network traffic is encrypted
- Non-business use of systems containing PII prohibited
- ... and the list goes on. You already KNOW what to do!

Next, Minimize PII Collected, Stored and Shared in the First Place

- *Q.: Who institutionally reviews and approves data to be collected, stored and shared at your college or university?*
- Avoid "high value" PII data such as social security numbers; use an institutionally assigned ID number instead (some use of SSNs may be unavoidable, e.g., for payroll purposes, and also perhaps in conjunction with insurance programs)
- Credit card data is subject to specific protection requirements described by Payment Card Industry Security Requirements (see https://sdp.mastercardintl.com/pdf/pcd_manual.pdf). Those PCI standards are not a bad starting for ALL systems with PII!
- Other critical data are defined by statute, including student records (FERPA), health records (HIPPA), & financial records (GLBA and/or voluntary institutional adoption of SarBox). 17

More "High Value" PII Worth Mentioning

- Critical records can be identified by their potential for causing institutional embarrassment or harm if disclosed, including:
 - hiring committee and promotion/tenure files
 - prospective donor assessments
 - confidential complaints & internal institutional legal opinions
 - proprietary/NDA'd commercial information
 - federal confidential/secret/top secret information
 - you can probably list many additional items here! :-)
- **Pay CLOSE attention to any document imaging projects!**
- Naturally, any/all PII record minimization effort needs to be consistent with institutional or state/federal data collection and record retention requirements, etc.
- Beware of institutional data shared with **off-site partners**; do you have liability for 3rd party breaches of shared PII?

When PII Is Stored on Disk or Backed Up to Tape or Disk, Ensure That It Is Encrypted

- Loss of **encrypted** PII data may not be considered “loss of PII” as defined under some statutes
- There may be offsetting institutional risk associated with potential **loss of access** to encrypted data that needs a now-lost or forgotten password/key (consider mandatory key escrow perhaps?)
- Key management can be a huge topic in and of itself! One example: how are keys stored and accessed/handled for automated processes (such as scheduled administrative "batch" job processing)?
- Speaking of keys, make sure keys do not get transmitted in the same package as backup media (for example, on yellow sticky notes taped to the outside of backup tape cases)

Encryption of Data on Mobile Devices

- Given the potential impact of lost laptops or other mobile devices, how should data be encrypted on **those** units?
 - manual encryption/decryption of individual sensitive files?
(may be inconvenient/forgotten/not routinely employed)
 - automatic decryption of files as needed, or automatic decryption/encryption of an entire file system "on the fly"?
(automated processes may potentially decrypt "too much," or allow access by an unauthorized user if an authorized user is working with PII when a bad guy/gal "hacks in")
- Some laptops come with integrated encryption support (e.g., Apple OS X offers File Vault), while in other cases you may need to add an external product such as TrueCrypt (a free open source laptop encryption project for Windows and Linux, see <http://www.truecrypt.org/>)

The Terminal Is Dead; Long Live the Terminal?

- Alternative approach: treat laptops as nothing more than a display device and store all sensitive files on a central server, requiring the remote user to login to the central server over a VPN or other encrypted link to access or modify PII data.
- Advantage? If the laptop's lost, there's no PII stored on it.
- Potential problems:
 - remote access to PII over the network will become routine (and so you trade elimination of laptop loss risk for a new potentially increased risk of unauthorized remote access)
 - think about/deal with/accept situations where users will be off-network (areas w/o coverage, time spent flying, etc.) or central systems go down (does that still happen? :-;)
 - application speed will be dependent on network quality
 - what about preventing copying of data to a local disk?
 - need to make sure locally created cache files get removed²¹

Theft of PII Data from Non-Mobile Devices (e.g., Desktop PCs)

Only a small number of the data breaches reported to the Committee were caused by hackers breaking into computer systems online. The vast majority of data losses arose from physical thefts of portable computers, drives, and disks, or unauthorized use of data by employees.

"Agency Data Breaches Since January 1, 2003, Committee on Government Reform, US House of Representatives," Oct 13, 2006
<http://www.govexec.com/pdfs/AgencyBreachSummaryFinal.doc>

- It isn't hard to crack the case on a typical desktop PC and steal the hard drive (so why aren't desktop PCs shielded in cradles, or why don't institutions buy desktop PCs with removable hard drives which can be locked in a safes when the PC isn't in use? Or should desktop PC hard drives simply be encrypted the same way that laptop hard drives are?) 22

Stealing Information Doesn't Require Physical Asset Removal...

- There's no need to physically remove part of a PC if we can just copy the information on the PC to removable media....
- Most PCs have one or more ways to export data w/o using the network, e.g., floppy disk drives, CD or DVD writers, USB ports, locally attached printers.
- Should those options be available on devices that handle PII data, or should PCs handling PII be stripped of those options? (Yes, I know this sounds really draconian)
- USB ports may be the hardest to control: a USB port left live for use in conjunction with a USB keyboard may end up being used to connect a USB hub, thereby allowing the USB keyboard AND a USB thumb drive to BOTH be connected....

PII and Logging

- Strive to do extensive logging to a non-modifiable device (remember paper console logs on numbered printouts?); beware of the possibility of logs on-machine being sanitized by an attacker, or network-based logs being blocked
- Accurate timestamps matter; sync system clocks with NTP!
- Be sure you can *promptly* detect signs of a system intrusion via system monitoring tools, and be sure you can also detect what files may have been modified via Tripwire, etc.
- The longer an intrusion or other incident goes undetected, the greater the probability that PII-related information "jewels" will be stumbled upon and exported
- The longer an incident goes undetected, the greater the popular perception that the "Captain was asleep at the wheel of the ship" (and thus the greater the popular outrage)

PII Data Inside Log Files & Status Pages

- As you look at PII data exposure, there's one area that's all too often overlooked, and that's PII data **INSIDE** public log files. Log files can be of critical importance when it comes to debugging system issues and attributing potentially malicious behavior, and for statistical summarization purposes, but raw log files can also publicly expose a tremendous amount of information about communication patterns, etc.
- Log files are often routinely publicly readable by any user on the system, and it is rare for there to be any compartmentalization of log data (e.g., I can see my data in the log file, but I can also see everyone else's data, too).
- Be sure to also consider things like dynamic server status pages (try doing <http://www.whatever.edu/server-status>)
- If you're not paying attention to PII data **INSIDE** log files and in dynamic status pages, you really, really, should be.

Passive Intrusion Detection Systems

- Beyond logging, intrusion detection systems (such as Snort or Bro) can be very helpful in timely detection of an intrusion, and passive monitoring can also lay the foundation needed for network forensics – e.g., does it look like the miscreants actually copied any files off the host they compromised? If so, where did those files get copied to?
- Passive network monitoring hosts may also represent a new PII breach vector because they see ALL traffic, so passive network monitoring hosts need to be carefully secured and controlled (network traffic for passive monitoring purposes should be delivered via unidirectional optical taps to an otherwise-off-net host for analysis).
- Be sure users get informed that an intrusion detection system doing passive network monitoring has been installed.
- What about access to systems by authorized users.

Authorized Users ==> “Accounts”

- Be sure to understand how/when accounts get created, and how accounts get disabled or deleted (what triggers creation or deletion?) Can you verify eligibility/need for each account?
- To what systems does a username/password give access?
- How do initial passwords get distributed? Are periodic changes required? Are passwords tested for crackability?
- How do passwords get reset when forgotten? What do they get reset to? How do you verify the identity of the requestor?
- Is access to an account by a third party (such as a coworker or supervisor) possible? If so, under what circumstances?
- How do usernames get assigned to user groups?
- What are the default protections on files in accounts?
- How are privileged accounts created, controlled and monitored? Are direct root logins allowed, or only sudo access? Are direct root logins from the network disallowed?

Account Management (continued)

- Beware of linkages or "chaining" that may occur: breach of system A results in theft of password hashes; cracking of those password hashes with Rainbow Table techniques (see, for example <http://www.antsight.com/zsl/rainbowcrack/>) yields username/password pairs that enable subsequent breach of systems B, C, D and E due to shared credentials/common username/password reuse.
- Single sign on is a common objective at many institutions but can undermine "natural firebreaks" that might have otherwise existed between groups of systems.
- Can you operationally issue and distribute all new passwords for thousands of accounts if circumstances require? (This can be a trickier issue than you might expect, particularly if you rely on email for many routine administrative communications, or you have lots of off campus users)

Risk Assessment

- Another part of managing PII exposure is knowing where PII resides in your institution (e.g., you need to do a risk assessment/vulnerability analysis)
- Many caches of PII live in central systems, but many others live in distributed/decentralized departmental systems, where extracts from central data sources have been downloaded and saved, and on personal workstations.
- If you don't know that PII data exists on a departmental system or on a personal workstation, you won't:
 - know that system needs to be checked for vulnerabilities and hardened, nor will you
 - know that that system merits special attention in the event of an incident because of PII-related considerations
- Recognize that some users may be reluctant to concede that PII even exists on distributed/decentralized systems

Data Classification

- Risk assessment/vulnerability analysis is facilitated if your institution has a formal system for data classification.
- Higher ed data classification systems often have three levels:
 - unrestricted/publicly available information
 - data that's non-public as a matter of institutional discretion
 - data protected from disclosure by law
- When picking labels for your classification levels, avoid any formal federally-defined document classifications such as confidential, secret, restricted, etc. (to avoid any confusion).
- Recognize that deciding "completely unlabeled == unrestricted and publicly available" has potential risks... what about a highly sensitive document that somehow escaped proper labeling, but which is definitely non-public?
- Don't forget about FOIA/open record/government in the sunshine laws if they apply to your record systems.

Penetration Testing

- Once you (think you) know where PII may be located on campus, you need to know if that PII is being stored securely.
- One way of checking this is by having a penetration test done, thereby possibly identifying vulnerabilities which can be corrected.
- Note: aggressive penetration testing (and that's the only worthwhile kind!) may result in mission critical systems going off line under some circumstances.
- Also note that once you're officially aware of a documented problem, it becomes much hard to ignore that problem, and regrettably, fixing some problems may be expensive.
- One last thought about penetration testing: pen testing can't formally and exhaustively prove that you ARE secure, it can only (potentially) provide evidence that you are NOT.

Policies and Procedures; Training; Build A Relationship with the Media

- Institutional policies and procedures should be in place supporting all of the above.
- University senior management and university legal counsel should review & affirmatively approve PII-impacting policies
- Technical measures and policies and procedures are all for naught if users don't know and embrace required activities. Users can't/won't do what they should unless they know what they should be doing. Thus, clearly, there also needs to be an ongoing user education/training initiative.
- Build a relationship with the local media. They can help you get the word out to users about what people should do to help minimize risks of PII exposure, and building a relationship with them now may pay off if/when a PII breach occurs at some later point.

Planning for the Worst

- While you hope you'll never experience a breach, you should still plan for the worst. If a breach does occur:
 - Who needs to be notified internally?
 - How will you notify potentially affected users?
 - Have you thought about information for the press?
 - Do you have procedures in place to obtain credit monitoring services for those who may be affected (at institutional expense)?
 - Can you restore potentially compromised systems or accounts from known-clean backups?
 - Do you have procedures for invalidating current passwords and securely issuing new ones?
 - Have you thought through whether you'll want to work with law enforcement to try to apprehend the perpetrator?
 - Do you have predetermined criteria for determining if a PII breach actually even occurred?

"Maybes" Can Be As Bad As "Dids"

- In some circumstances, when a system has been compromised and personally identifiable information may have been exposed, you may simply not be able to definitively tell if PII on that system actually was compromised or not... This is a VERY common scenario.
- "Maybe the bad guys got at some PII" can be as bad or worse than "the bad guys absolutely DID get at some PII..."
- Early detection along with network forensics may help to eliminate uncertainty for some compromises, but sometimes you simply may never know for sure.
- You can be careful and assume the worst, but that can be expensive and embarrassing; alternatively, you can be optimistic and hope for the best, but that may not be warranted and may cause its own problems. Carefully think through how you'll want to handle "maybes" in advance... 34

Cyberinsurance

- One final option to consider when it comes to mitigating the risk of PII spills is the purchase of cyberinsurance. A nice overview of this can be found in "Worried About Hackers? Buy Some Insurance," Chronicle of Higher Education, October 13, 2006, <http://chronicle.com/weekly/v53/i08/08a04101.htm>
- Cyberinsurance shouldn't be viewed as a "magic bullet:"
 - If you have 10,000 users and buy a \$3 million dollar policy, that only provides you with \$300 worth of coverage/user... enough to cover some PII notification and mitigation costs, but not enough to satisfy large scale lawsuits
 - You may not qualify for coverage. Insurers will commonly review your IT security policies and practices, and if you appear particularly vulnerable, coverage may be declined
 - There will often be material exclusions to your coverage₃₅

Monster #2: Malware (Viruses, Worms, Trojan Horses, Spyware, Root Kits, etc.)

"Frequency of attacks. Nearly nine out of 10 organizations experienced computer security incidents in a year's time; 20% of them indicated they had experienced 20 or more attacks.

"Types of attacks. Viruses (83.7%) and spyware (79.5%) headed the list..."

New FBI Computer Crime Survey, 1/18/06

www.fbi.gov/page2/jan06/computer_crime_survey011806.htm

The Most Frequent Monster

- We hear an awful lot about PII, but the real "worst" security monster, or at least the most frequently encountered IT security monster, is malware (viruses, worms, trojan horses, spyware, root kits, crimeware, etc.).
- Malware is a very real and ongoing problem for higher education.

Higher Education Desktop/Laptop Antivirus Licensing Practices

- Virtually every college and university in the country has licensed a desktop/laptop antivirus (AV) product for at least some parts of their campus community.
- Unfortunately many times that coverage is incomplete:
 - the college licensed an AV product but not antispyware
 - the college licensed an AV product for use by faculty and staff, but not for students (or vice versa, often due to the funding of the AV license by student ed tech fees)
 - the college licensed a product for on campus use, but didn't also purchase coverage for home systems
 - users may have the product installed, but for one reason or another updated antivirus definitions aren't getting downloaded and installed
- There can be other AV-related issues, too...

Example: AV Signatures

- Most AV products are "signature based," and identify viruses based on peculiarities ("signatures") unique to each virus.
- New virus signatures only get released by the vendor and downloaded by the end user perhaps once a day, while miscreants can release new not-yet-detectable versions of their malware as often as they want (e.g., multiple times a day). The virus writer can thus guarantee that they will have a period of time during which user systems will be vulnerable.
- Virus writers also enjoy another key advantage: they can empirically test and repeatedly tweak their code and its packaging until their exploit doesn't get detected by current popular antivirus products.
- Thus, it is a virtual certainty that at least some malware will get pass your current AV solution... But most users don't understand that.

AV Products Can Give Users A False Sense of Security

- So beware the fearless "user warrior" who bravely roams at will online, confident that his antivirus software will shield him from any malware he might run accidentally run into...
Antivirus products can help reduce the risk of an infection, but they don't, and can't, grant comprehensive immunity.
- Users may also believe that if they do somehow get infected, their antivirus product can act as a magic "cyber antibiotic," and successfully clean up any infection they've managed to acquire, leaving their infected system good as new after the AV gets done running. Of course, security professionals know that that will often not be true, and the only sure way to get a clean and stable system again is to "nuke and pave" the system, reinstalling the OS and all applications from scratch, and restoring user files from known clean backups.
- But do we make sure our users know these sort of things?⁴⁰

Testing AV Coverage

- An interesting exercise if you're looking for a new geek passtime: carefully submit any suspicious executables you come across which aren't flagged by your current antivirus program to VirusTotal and/or to Jotti, and see how many of them end up getting detected by any of the dozen or more popular antivirus software products those sites use...
<http://www.virustotal.com/>
<http://virusscan.jotti.org/>
- Bummer: MANY suspicious files will turn out to be malicious, and will get flagged by one or more AV products, but at the same time that malware will be missed (or misidentified) by many of the other AV products running on those sites.
- Want to get REALLY bummed out? Check the same file again days later; notice how often a given piece of malware STILL doesn't get detected, even though Virustotal and Jotti share submitted executables with participating AV vendors.⁴¹

Periodically "Double Checking" Things

- Once you recognize that antivirus coverage may be incomplete at best, you may want to get a "second opinion" when it comes to the cleanliness of your system, just in case something "slid by" your normal antivirus product. A couple examples of free online scanning tools include:

- <http://usa.kaspersky.com/services/free-virus-scanner.php>
- <http://housecall.trendmicro.com/>

Note that many of these type of tools use ActiveX, which means you'll be running them from Internet Explorer

- Another tool you should know about is MyNetWatchman's SecCheck, see <http://www.mynetwatchman.com/tools/sc/>
SecCheck does a very nice forensic review of an infected PC

A Matter of Semantics (and Marketing)

- In some cases, malware may be known to your AV company, it just doesn't get detected by your AV product because your AV company has decided to categorize that particular malware as "spyware" rather than a virus or trojan horse, and you've only licensed the company's AV product (and not also the company's antispymware product).
- Don't get hung up in an argument over semantics: you want to detect and block as much malicious content as possible, regardless of what it is called or how a vendor categorizes it or how it propagates onto a user's system.
- License both the antivirus and the antispymware product from whatever vendor you select!

The Problem of Trial AV Coverage

- Another AV marketing-related issue: AV companies are sure being nice when they offer free ninety day trial coverage as part of a new system software bundle, right? Wrong.
- Frequently users receive a free trial antivirus product as part of a new system software bundle, but then, when the free trial period runs out, they fail to buy the product or subscribe to get continued antivirus signature updates. Naturally an AV product without signature updates offers pretty incomplete protection – although many non-technical users will overlook this ("Hey, my AV product is still running, right?")
- AVOID short term/trial AV coverage to the maximum extent possible. Be sure users don't incorrectly choose to use the (temporarily) "free" antivirus product that comes with their new system instead of the site licensed antivirus product you've provide for their use!

Defense In Depth: AV on Servers

- You can improve your chances of blocking malware by using multiple AV products, one on your campus desktops/laptops, and another product from a different vendor on your servers. For example, UO uses ClamAV on our central servers, and McAfee on our desktops/laptops, thereby increasing the chance that a virus missed by one AV product may get detected and handled by the other. (Note that ClamAV is a free product, so lack of budget is no excuse when it comes to running a true antivirus product on your servers!)
- UO also uses Procmail Email Sanitizer (PES) to defang or strip potentially dangerous content that's being sent by email, taking action based on the file extension of the attachment involved (thereby avoiding at least part of the problem with the AV good guy/bad guy signature "race"). For info on PES, see <http://www.impsec.org/email-tools/procmail-security.html>
- PES also defangs risky HTML elements from incoming mail

The Problem of HTML Formatted Email

- The fundamental issue:
 - increasing numbers of users are sending HTML-formatted email by default (e.g., the traditional send "plain text email only" Internet culture is rapidly being exterminated)
 - HTML-formatted email can be exploited to download or run malicious content in an amazing number of different ways (see some examples at <http://ha.ckers.org/xss.html>)
 - the vast majority of users run with scripting enabled (see http://www.w3schools.com/browsers/browsers_stats.asp)
 - rendering HTML-formatted mail safe(r) typically breaks the HTML formatting of most HTML-formatted messages
 - if you don't sanitize HTML-formatted mail, your users WILL get 0wn3d
 - if you do sanitize HTML-formatted mail users WILL complain
- Try to emphasize/require plain text email whenever possible

OS and Application Choice

- **Observation:** Virtually all currently known malware targets systems which run Microsoft Windows.
- **By implication:** one of the simplest things you can do to avoid problems with malware is to NOT run Microsoft Windows. You **do** have options!
- **Observation:** Different mainstream applications DO have different risk profiles. Do you know how many vulnerabilities have been reported for the applications you use? Looking at just unpatched vulnerabilities -- how many remain unpatched? What's the highest severity vulnerability associated with a currently unpatched vulnerability?
- **One excellent source for that sort of data:** <http://secunia.com>
Note: we're about to do precisely the sort of head-to-head comparison that Secunia officially explicitly discourages...⁴⁷

An Example: Operating Systems

Checking Secunia on October 30th, 2006...

- **Windows XP Pro (<http://secunia.com/product/22/>):**
157 advisories
28 unpatched
most severe unpatched: highly critical
- Apple Mac OS X (<http://secunia.com/product/96/>):
75 advisories
0 unpatched
- Red Hat Fedora Core 5 (<http://secunia.com/product/8808/>):
10 advisories
0 unpatched

Another Example: Web Browsers

- **Internet Explorer 6.x** (<http://secunia.com/product/11/>):
106 advisories
19 unpatched
most severe unpatched: highly critical
- **Internet Explorer 7.x** (<http://secunia.com/product/12366/>):
3 advisories,
3 unpatched
most severe unpatched: moderately critical
- **Mozilla Firefox 1.7.x** (<http://secunia.com/product/3691/>):
36 advisories
6 unpatched (Firefox 2.x currently has 0 advisories)
most severe unpatched: less critical
- **Opera 9.x** (<http://secunia.com/product/10615/>):
2 advisory
0 unpatched

Email Clients...

- **MS Outlook Express 6 (<http://secunia.com/product/102/>):**
22 advisories
7 unpatched
most severe unpatched: moderately critical
- Mozilla Thunderbird 1.5.x (<http://secunia.com/product/4652/>):
4 advisories
0 unpatched
- Recognize that many casual users will just use a web email interface; look for one that will allow them to work with their email without having to have Javascript enabled in their browser. One example to consider is UO's free/open source web email product: <http://alphamail.uoosl.org/>

I Know It May Not Be Easy...

- I know it may not be easy to swim against the tide. For example, your ERP product (or your bank, or your favorite web shopping sites) may only work right with certain browsers, or your institution may standardize on a single messaging and calendaring client (even if it has known vulnerabilities).
- If that happens, you may want to:
 - try to educate local decision makers about the risks,
 - make sure vendors and web sites know how important it is for them to support all standards compliant browsers,
 - use safer operating systems and applications when you do have the option, using less safe options only when you absolutely have no other choice.

Applications Other Than Web & Email

- While the web and email are two particularly critical applications, depending on your local institutional culture, other applications may also be quite important, including things like P2P file sharing applications, Usenet News, instant messaging or IRC, dedicated RSS clients, etc.
- Be sure to also pay attention to external helper applications and plugins (Acrobat Reader, Java, Quicktime, Real, etc.).
- One specific example of an external helper-related issue: installing a new version of Java does not automatically remove any old (vulnerable) versions which may also be installed, and that installation behavior is intentional, not a bug. Nice discussion of this issue by Brian Krebs at "Sun Acknowledges Security Hole In Patch Process," http://blog.washingtonpost.com/securityfix/2006/08/sun_acknowledges_major_oops_in.html

Rootkits

- You scan for viruses and spyware, but what about rootkits?
- Rootkits help malware hide on systems, just as stealth technologies help aircraft evade detection by radar.
- One rootkit detector is RootkitRevealer by Sysinternals; see: <http://www.sysinternals.com/Utilities/RootkitRevealer.html>
Unfortunately, unlike many antivirus and antispymware applications, RootkitRevealer really isn't a suitable tool for non-technical users (it is more of a specialist's tool), requiring some expertise to interpret its output.
- Some other rootkit detection products to try are:
 - F-Secure BlackLight (http://www.f-secure.com/blacklight/try_blacklight.html), currently in free beta thru 1 Jan 2007
 - Sophos Anti-RootKit (<http://www.sophos.com/products/free-tools/sophos-anti-rootkit/download/>)

Monster #3: Distributed Denial of Service Attacks, Spoofed Traffic, BCP 38 Filtering and Open Recursive Name Server Abuse

"More than five years after the initial flurry of network attacks, and the news articles and research papers that followed, DDoS remains the number one concern for large IP network operators.

Sixty-four percent of the survey participants said, 'DDoS is the most significant operational security issue we face today.'"

Worldwide ISP Security Report, September 2005

[http://www.arbornetworks.com/downloads/
Arbor_Worldwide_ISP_Security_Report.pdf](http://www.arbornetworks.com/downloads/Arbor_Worldwide_ISP_Security_Report.pdf)

The Monster That Can Bite The Hardest

- As troubling as PII breaches can be, and as ubiquitous as malware can be, the IT "monster" that unquestionably can "bite the hardest" and "hurt the most" is the distributed denial of service (DDoS) attack.
- But what is a DDoS attack?

Examples of DDoS Attacks

- In a distributed denial of service attack, network traffic from thousands of hacked computer systems -- often systems located all over the Internet -- gets used in a coordinated way to overwhelm a targeted network or computer, thereby preventing it from doing its normal work. For example:
 - the institution's connection or connections to the Internet may be made to overflow with unsolicited traffic (a so-called "packet flood")
 - web servers may be inundated with malicious repeated requests for web pages
 - campus name servers may become swamped so that university computer users have problems visiting either local web sites or web sites on the Internet

Effects of a DDoS:

- The systems and networks that are the target of the DDoS attacks are still there and haven't been hacked or compromised, BUT they are too crippled to do useful work.
- An attack that is targeting a single server or desktop can have collateral damage on an entire site to the extent that infrastructure (such as a common Internet connection) is shared.
- When the denial of service attack stops or is abated, the targeted systems are usually able to rapidly resume normal operation; lingering effects should be minimal or non-existent.
- Blocking or abating one DDoS usually will not prevent another from occurring.

So What's The Big Deal? Why Not Just Filter The Problematic Traffic?

- It can be tricky to filter the attack traffic.
- For example, if your connection is being flooded with inbound traffic, you need to block it upstream, BEFORE it can traverse the last network links into your site. If you try to filter the traffic at your campus border it will be too late at that point – your inbound network pipe will still be unusably full.
- The miscreant DDoS'ing you may have an army of tens of thousands (or hundreds of thousands of compromised hosts)... and the hosts he's using may constantly change.
- Attackers may change their attack mechanism over time, adapting to blocks you put into place.
- There are some types of attacks where it is extremely hard to characterize attack traffic in a way that will allow it to be distinguished from legitimate traffic in a filterable way.

How About This: What If We Treat It Like A Blizzard, And Just 'Ride It Out?'"

- While there is a certain insouciance to the idea of having "denial-of-service days" (sort of like more traditional "snow days"), higher education folks should understand that denial of service attacks can be sustained for days -- or even weeks or more -- at a time. For example, Spamhaus, a major anti-spam activist organization, was subject to an attempted denial of service attack that lasted for **three months**. (See <http://www.spamhaus.org/news.lasso?article=13>)

Taking an entire denial-of-service term off would have material impacts on a university's ongoing operations, and probably would simply be unacceptable.

Let's Just Disconnect For a While

- While disconnecting from the Internet would certainly insure that attack traffic coming from the Internet cannot DDoS university systems, disconnecting entirely from the Internet is itself a form of self-imposed denial of service, and would likely not be well received by campus constituents.
- In the case of inbound DDoS attacks targeting a particular non-mission critical host, disconnecting that single host may be a pragmatically viable strategy...
- Likewise, in cases where a single compromised host is being used to generate outbound flows, disconnecting that compromised host is almost always the right thing to do (unless you're trying to collect forensic evidence from that live compromised host for prosecution).
- Speaking of law enforcement...

Call the FBI and Let Them Sort It All Out

- The FBI and other law enforcement officials will typically be interested in major DDoS attacks, however their attention will not provide symptomatic relief when a DoS occurs, nor is it a guarantee of a successful investigation and eventual prosecution – DDoS cases can be hard to put together.
- You should also understand that many times denial of service attacks are transnational, which introduces special investigatory issues, and requires FBI coordination with foreign LE counterparts, which can introduce substantial investigative delays. Denial of service attacks committed by individuals overseas (and attacks made by minors whether here in the US or abroad), may also result in disappointingly short sentences. This may dampen LE/DA enthusiasm for proceeding with a potentially hard-to-investigate, hard-to-prosecute, low-payoff case.

Is Higher Education An Attractive Target For A DDoS-Based Extortion Attempt?

- Imagine a threatened DDoS attack during a crucial time, such as during a prime window for students to submit applications for admission – how many of us now rely on online applications for a significant proportion of our matriculating class? How tight is that window? Do you routinely send out printed backup application materials?
- Or maybe you have closely defined windows for students to enroll in classes via an online portal -- what would the impact be if your enrollment system was offline for half a day or a day during peak registration times? Or how long could you continue to function without access to your institutional teaching and learning system? Or your administrative ERP system?
- I think higher education IS vulnerable to DDoS extortion. 62

DDoS Identification

- One of the hardest problems you may initially face if you do get hit is simply identifying that a DDoS is going on....
- Some institutions may not have formal network monitoring in place, and so a result the first indication that "something's wrong" may be user complaints.
- Once your staff begins to suspect that something is wrong, differential diagnosis will require them to first rule out the possibility that systems are just experiencing higher-than-normal "real" loads.
- The next issue then becomes where's the load? Is the load on the network? On a single server? On a set of servers? On some piece of supporting infrastructure like campus name servers or web cache servers?
- Can we tell when the DDoS started? Is it still going on?

DDoS Identification Can Be One of the Easiest Things to Tackle

- If you systematically work to improve your network monitoring, identifying the fact that a denial of service attack is occurring will quickly become routine.
- Do you have MRTG or RRDtool graphs that monitor your network traffic levels in octets and packets per second? (strip charts of that sort are extremely helpful when it comes to tracking macroscopic DDoS behaviors in a management-friendly way).
- Do you have an intrusion detection system, such as a Snort or Bro box deployed? If not, it will be an excellent investment.

DDoS Mitigation

- Mitigating a distributed denial of service attack is usually a collaborative process, and will usually involve you or your institution's networking staff working on the phone with ISP's network engineers and security staff, etc.
- Do your networking staff know your ISP's engineers and security staff? If not, this might be something to work on rectifying BEFORE a denial of service attack occurs. Personal relationships can/do matter when it comes to mitigating denial of service attacks!
- Depending on the ISP you use, you actually may have a more efficient technical option available to you, known as "blackhole communities."

Directly Sinking Attack Traffic Via Blackhole Communities

- If you're fortunate, your ISP may allow downstream customers to self-tag routes with blackhole community values following the process outlined at <http://www.secsup.org/CustomerBlackHole/> or as discussed in more detail at <http://www.nanog.org/mtg-0410/pdf/soricelli.pdf>
This approach allows attack traffic to be blackholed by a targeted site in an efficient fashion, as close to the attack source as possible.
- **Suggested Investigation Item For You or Your Staff:**
Does your ISP support blackhole communities?
If so, do you know what values to use if you need them?

Learn More About DDoS Before You Get Hit

- Some excellent technical papers include:
 - Hank Nusbacher's "DDoS: Undeniably a Global Internet Problem Looking for a Global Solution,"
<http://www.interall.co.il/presentations/ripe41.pdf>
 - HoneyNet's "Know Your Enemy: Tracking Botnets"
<http://www.honeynet.org/papers/bots/>
 - John Kristoff's "Botnets" talk from NANOG 32
<http://aharp.ittns.northwestern.edu/slides/botnets.pdf>
 - Peter Moody's Botnets talk from the SLC Joint Techs
<http://www.internet2.edu/presentations/jtsaltlake/20050214-Botnets-Moody.pdf>
 - More resources:
<http://www.honeypots.net/incidents/ddos-mitigation>

The Role of Spoofed Traffic in DDoS Attacks

- Now that you understand some of the implications of a DDoS attack, we can move on to talking about how DDoS attacks often end up being implemented.
- A key DDoS-enabling technology is spoofed traffic.
- To understand how spoofed traffic works, we first need to talk a little about the types of network traffic we see on the network.

TCP and UDP Traffic

- There are basically two types of network application traffic: TCP and UDP.
- TCP traffic is associated with relatively persistent connections (such as ssh sessions, web traffic, email, etc.), and has a variety of characteristics which are desirable from a network application programmer's point of view, including retransmission of lost packets, congestion control, etc.
- UDP traffic, on the other hand, is designed for "send-it-and-forget-it" applications where you don't want to/can't afford to maintain state or you don't want a lot of connection setup overhead. DNS, NFS, and IP video traffic all normally run as UDP.

The Spoofability of UDP Connections

- Unlike a fully established TCP connection (which only gets established after a bidirectional handshake is negotiated and which is therefore robust to spoofing attempts),* UDP traffic can be created with virtually **any** apparent source address -- including IP addresses which have no relationship to the traffic's actual origin.
- Network traffic that's intentionally created with a bogus source address is said to be "spoofed."
- If allowed to reach the global Internet, spoofed traffic is generally indistinguishable from legitimate traffic.

* Yes, of course, naked TCP SYNs are also spoofable.

Why Would Anyone Spoof Traffic?

- If you don't spend time "thinking like an attacker," you might not immediately "get" why an attacker would be interested in spoofing his attack traffic. The answer is actually quite simple: the attacker wants the systems he's using as part of his attack to stay online and unblocked as long as possible.
- Spoofing the source of the attack traffic...
 - hinders backtracking/identification/cleanup of the system that's sourcing the traffic; and
 - makes it harder for the attack victim to filter the attack traffic (the spoofed source addresses may be constantly changed by the attacker, and thus they won't provide a stable "filterable characteristic").

So Why Not Just Block All UDP Traffic?

- Given that UDP can be easily spoofed by the bad guys/bad gals, sometimes you'll hear folks naively propose simply blocking all inbound or outbound UDP traffic (or at least heavily rate limiting all UDP traffic).
- Unfortunately, because some pretty basic services (like DNS) requires support for UDP, blocking (or heavily rate limiting) all inbound or outbound UDP traffic is generally **not** a good idea. Warts and all, you have no choice but to learn to to live with UDP traffic. :-;

Well, Can We Block SOME UDP Traffic?

- For once, the answer is positive: yes, you can block some UDP traffic.
- For example, if you're the University of Oregon and your school has been assigned the IP address range 128.223.0.0-128.223.255.255 there's no reason for systems on your network to be sourcing packets that pretend to be from some other IP address range. We'd filter that spoofed traffic before it leaves our campus.
- This is a pretty basic sanity check, but you'd be surprised how many sites don't bother with even this trivial sort of filter.

Subnet-Level Filtering

- While it is great to prevent spoofing at the university-wide level, that sort of border router anti-spoofing filter does not prevent a miscreant from forging an IP address taken from one of your subnets for use on another of your subnets.
- *Cue subnet-level anti-spoofing filters....*

You KNOW that hosts on each subnet should ONLY be originating packets with IP addresses legitimately assigned to that subnet, so at the uplink from each subnet, drop/block outbound packets that appear to be "from" any other IP address – another very basic sanity check.

BCP38/RFC2827

- Let me be clear: ingress filtering of traffic with spoofed IP addresses is not new and is not my idea – it is Best Current Practice (BCP) 38/RFC 2827,* written by Ferguson and Senie in May 2000.
- Unfortunately, despite being roughly six years old, **many** sites still do **NOT** do BCP38 filtering -- perhaps as many as 20-25% Internet wide.**
- **Does YOUR school do BCP38 filtering? If not, you really should!**

* <http://www.ietf.org/rfc2827.txt>

** <http://spoofer.csail.mit.edu/summary.php>

So Why Doesn't Everyone Do BCP38 Filtering?

- Asymmetric costs/benefits: filtering my network protects you (which is nice), but filtering that traffic "costs" me w/o any tangible/economic "benefits." What are these "costs?"
 - engineer time to configure and maintain the filters (one time/negligible for most .edu networks)
 - overhead on the routers (but if that overhead is material enough to be a "show stopper," you should be upgrading anyway)
- Other common (lame) excuses:
 - "Too hard given the complexity of my network"
 - "I'm too busy"

What's It To You Anyway, Bub?

- Some may question why others should care what they do with their networks – their network, their rules, right? Well, generally yes. However in this case, remember that if someone's NOT doing BCP38 filtering, that network may be getting used to generate spoofed attack traffic that's pretending to be "from" an innocent third party network, and that's the point at which what someone does (or doesn't do) potentially affects a lot of other people including the attack target itself, the entity whose IP addresses are being spoofed, etc.
- It's important to be a good neighbor.
- Make **sure** you're doing BCP 38 filtering of spoofed traffic!

So How Should I Be Doing BCP38 Filtering?

- Only you (and your network engineering team) can make the final decision about the best approach for your network, but you may want to see BCP84/RFC3704, March 2004.
- I would note, however, that strict mode unicast reverse path forwarding ("strict uRPF") may not be a good idea for the multihomed environment typical of many universities due to route asymmetry. I would also urge you to review draft-savola-bcp84-urpf-experiences-00.txt (April 19, 2006)
- One answer, quoting from RFC3704 (mentioned above):
"Ingress Access Lists require typically manual maintenance, but are the most bulletproof when done properly..."

A Specific Example of UDP Spoofing

- Since we just got done covering UDP spoofing, talking a little about open recursive domain name servers and DNS amplification attacks seems like a "nice" segue/practical example of why BCP38 filtering is important, while also pointing out another specific vulnerability you should be addressing.
- Again, let's begin with a little background, first.

Thinking About DNS

- Most users never really think about how DNS works* -- they just take it for granted that entering `http://www.uoregon.edu/` in their web browser will take them to the University of Oregon home page.
- In order for that to happen, however, the web browser needs to be able to find out that `www.uoregon.edu` resolves to the IP address (or "dotted quad") `128.223.142.13`
- The web browser, and ultimately the user, relies on the domain name system to do that name-to-dotted quad translation.
- DNS is thus a critical network service.

* Geeks may see RFC1035 for the gory details.

Authoritative & Recursive DNS Servers

- There are different types of domain name servers, with "authoritative" and "recursive" DNS servers being the two most important types:
 - Authoritative servers are definitive for particular domains, and should provides information about those domains (and ONLY those domains) to anyone.
 - Recursive servers are customer-facing name servers that should answer DNS queries for customers (and ONLY for customers) concerning any domain.
- DNS servers that aren't appropriately limited can become abused.

For Example...

- Consider a situation where a DNS server is recursive AND is open for use by anyone (a server that's cleverly termed an "open recursive DNS server").
- While it might seem sort of "neighborly" to share your name server with others, in fact it is a really bad idea (the domain name system equivalent of running an open/abusable SMTP relay, in fact).
- The problem? Well, there are actually **multiple** problems, but one of the most important ones is associated with spoofed UDP traffic (see how this all ties together? :-;)

Spoofer DNS Attack Scenario

Dramatis personae:

- Attacker, who's working from non-BCP38 filtered network. Let's call him/her "A"
- Attack target – let's refer to that entity as "T"
- Open recursive domain name server on large, high bandwidth pipe, denoted below as "NS"

Act 1, Scene 1:

- "A" generates spoofed DNS queries, pretending to be "T," with "T"'s address as the "source" address of those queries
- "NS" receives the spoofed queries and dutifully returns the "responses" for those queries to "T" (even though "T" actually didn't ask for that information)
- "A" repeatedly generates such queries, DoS'ing "T" via "NS"

Spoofer DNS Attack Scenario Notes

- From "T"'s point of view, the attack comes from "NS," not from "A"
- DNS queries are small and use UDP, so an attacker can generate a "large" query volume
- DNS response traffic is also UDP, which means that it is insensitive to net congestion, and doesn't back off
- **DNS responses can be large relative to size of DNS queries (output/input ratios can run 8X+ on most DNS servers, and on servers supporting RFC2671 EDNS0 extensions, observed amplification can be >70X).**
- "A" can employ **multiple query sources (to generate a higher volume of spoofed queries), and use multiple NS's** to scale up the volume of response traffic generated (oh boy!)
- **Multi gigabit per second attacks can be generated.**

This Is A Well Known Vulnerability

- I'm not letting the "cat out of the bag" about a big secret -- this is a well known/documented threat. See, for example:
 - "The Continuing Denial of Service Threat Posed by DNS Recursion"*
 - "DNS Amplification Attacks"**
 - "DNS Distributed Denial of Service (DDoS) Attacks"***

* http://www.us-cert.gov/reading_room/DNS-recursion121605.pdf

** <http://www.isotf.org/news/DNS-Amplification-Attacks.pdf>

*** <http://www.icann.org/committees/security/dns-ddos-advisory-31mar06.pdf>

Open DNS Servers Worldwide

- Unfortunately, despite this being a well known problem, it is estimated that 75% of all name servers worldwide run as open recursive name servers.*
- And in a spirit of self-criticism, feel free to note that UO's name servers were open until we secured them this past February 1st, 2006.**
- **If *our* domain name servers were open recursive until Feb 2006, *how about yours?* You can check [your domain](http://dnsreport.com/) at <http://dnsreport.com/>**
- **If your DNS servers appear to be open recursive, you NEED to get them secured ASAP.**

* <http://dns.measurement-factory.com/surveys/sum1.html>

** <http://cc.uoregon.edu/cnews/winter2006/recursive.htm>

The Problem Isn't "Just" About DDoS

- By the way, if you aren't yet sufficiently motivated to "bite the bullet" and fix your DDoS-exploitable domain name servers, let me add a little more thrust to help launch that hog: if you're not controlling access to your domain name servers, you may also be leaving yourself vulnerable to **DNS cache poisoning attacks**, whereby vulnerable caching name servers can be made to return bogus results for a user's name service queries.*

* <http://www.lurhq.com/dnscache.pdf>

What's a Cache Poisoning Attack?

- In a nutshell, in cache poisoning attacks, the attacker "primes" the caching name server to respond to queries with an IP address of his/her choice, rather than the real/normal IP address for that site.
- The innocent victim asks the caching name server for the IP address of a site of interest, such as the IP address of their bank's website.
- If the domain name of that site happens to be one that the attacker has poisoned, the victim is automatically and transparently misdirected to a website of the attacker's choice, rather than to their bank's real web site, and confidential data can then end up being lost.

Another Cache Poisoning Scenario

- Another cache poisoning scenario uses cache poisoning to redirect queries for popular sites (such as google.com or hotmail.com) to a site that contains a virus or other malware.
- If your caching name server has been poisoned, when you try to visit one of these popular sites, you can unknowingly be redirected to another site that stealthily tries to infect your PC with malware.
- Blocking open access to your recursive name servers won't completely eliminate the possibility of your servers participating in such attacks, but it will reduce the likelihood of that sort of abuse.

Some DNS Recommendations

- Insure that you're running a current version of BIND* (or whatever DNS software you're using)
- Insure that you've separated your Internet-facing authoritative DNS server from your customer-facing recursive DNS server.
- Protect your user-facing recursive name server from access by users from other sites
- Consider analyzing DNS traffic with DNStop**
- Consider donating DNS log data to the RUS-CERT Passive DNS Replication Project***

* <http://www.isc.org/index.pl?/sw/bind>

** <http://dns.measurement-factory.com/tools/dnstop/>

*** <http://cert.uni-stuttgart.de/stats/dns-replication.php>

Monster #4: Address Space Hijacking

The Sneakiest Monster

- Address space hijacking is a problem that you may not even have heard of – it is the sneakiest monster we'll talk about today.
- When someone engages in address space hijacking, they use a range of IP addresses without proper authorization, and as you know, online, IP addresses effectively ARE your identity.

A Connection From UO, Right?

- Assume you saw a connection to your server from the IP address 128.223.142.13
- If you checked the DNS for that address on a Unix box, or if you checked whois, you'd associate that address with UO:

```
% host 128.223.142.13
13.142.223.128.in-addr.arpa domain name pointer darkwing.uoregon.edu.
% host darkwing.uoregon.edu
darkwing.uoregon.edu has address 128.223.142.13
```

```
% whois -h whois.arin.net 128.223.142.13
OrgName:      University of Oregon
OrgID:        UNIVER-193
Address:      1225 Kincaid St
City:         Eugene
StateProv:    OR
PostalCode:   97403-1212
[etc]
```

In Reality, However...

- **Just because some IP address is shown as having been assigned or allocated to someone doesn't mean that they're the ones actually USING that address.**
- For example, a miscreant may be able to arrange to have a third party ISP announce ("route") a range of IP addresses which they don't legitimately control. That announcement can be persistent, or temporary (e.g., brought up just long enough for a spam run and then withdrawn), a processes commonly known as "address space hijacking."
- **Address space hijacking may have important implications for network security activities which rely on the backtracking of observed connections.**
- **If address space hijacking is occurring and you don't consider that possibility, you may end up "going after" the wrong (innocent) party.**

Even The Feds Are Focused on IP Usage and Attribution Information

- The belief that if you "know" an IP (and a timestamp/time zone in the case of dynamic addresses) you "should" be able to tell who's associated with that address is reflected in customer record retention requirements mentioned in:
 - The Attorney General's remarks at the NCMEC:
www.usdoj.gov/ag/speeches/2006/ag_speech_060420.html
 - Congresswoman Diana DeGettes's ISP data retention requirement amendment:
energycommerce.house.gov/108/Markups/04262006/degette_001_XML.PDF
 - EU/International data retention programs
www.epic.org/privacy/intl/data_retention.html
- **Policy makers need to understand that apparent Internet traffic sources can not be taken at face value; routing of IP addresses must also be carefully considered.**

A Hand-Waving Introduction to Routing

- A "route" can (informally) be thought of as the **path** that network traffic takes as it proceeds from its source to its destination. Anyone who's used the **traceroute** command has seen examples of network paths. For example, lets trace to 128.223.142.13 from a looking glass server in Seattle (for a list of looking glass sites see <http://www.traceroute.org>):

```
Tracing the route to darkwing.uoregon.edu (128.223.142.13)
 0  so-3-0-0.gar1.Seattle1.Level3.net (209.0.227.133)  0 ms  4 ms  0 ms
 1  ge-11-1.hsa2.Seattle1.Level3.net (4.68.105.103)  [AS 3356]  0 ms
    ge-10-2.hsa2.Seattle1.Level3.net (4.68.105.135)  [AS 3356]  0 ms
    ge-11-1.hsa2.Seattle1.Level3.net (4.68.105.103)  [AS 3356]  0 ms
 2  nero-gw.Level3.net (63.211.200.246)  [AS 3356]  12 ms  4 ms  4 ms
 3  ptck-core2-gw.nero.net (207.98.64.138)  [AS 3701]  4 ms  4 ms  4 ms
 4  eugn-core2-gw.nero.net (207.98.64.1)  [AS 3701]  8 ms  4 ms  8 ms
 5  eugn-car1-gw.nero.net (207.98.64.165)  [AS 3701]  8 ms  8 ms  8 ms
 6  uonet8-gw.nero.net (207.98.64.66)  [AS 3701]  4 ms  8 ms  4 ms
 7  ge-5-1.uonet2-gw.uoregon.edu (128.223.2.2)  [AS 3582]  8 ms  8 ms  8 ms
 8  darkwing.uoregon.edu (128.223.142.13)  [AS 3582]  8 ms  4 ms  8 ms
```

Looking At That Traceroute...

- That traceroute shows the hop-by-hop path that traffic took going from a host in Seattle to 128.223.142.13. Because that traceroute was done from a "looking glass" running on a router, besides showing us "normal" traceroute stuff (such dotted quads and the host names for each hop in the path), it **also** shows us some **additional** numbers, e.g.: **"AS 3356," "AS 3701,"** and **"AS 3582."**
- Those numbers represent the "autonomous systems" through which network traffic might pass when going from our source host to our destination host. **AS3356** represents Level3, **AS3701** represents NERO (Oregon's higher education network), and **AS3582** represents the U of O. That is a perfectly reasonable path for traffic to take in this case.
- The **last AS in the path** shows who's USING that IP, NOT just the entity to whom that IP address was assigned!

Traceroute From a Site in Switzerland

Tracing the route to darkwing.uoregon.edu (128.223.142.13)

```
1 switch.rt1.gen.ch.geant2.net (62.40.124.21) [AS 20965] 4 ms 0 ms 0 ms
2 so-7-2-0.rt1.fra.de.geant2.net (62.40.112.22) [AS 20965] 8 ms 8 ms 16 ms
3 abilene-wash-gw.rt1.fra.de.geant2.net (62.40.125.18) [AS 20965] 128 ms 124 ms
  112 ms
4 nycmng-washng.abilene.ucaid.edu (198.32.8.84) [AS 11537] 112 ms 108 ms 108 ms
5 chinng-nycmng.abilene.ucaid.edu (198.32.8.82) [AS 11537] 132 ms 132 ms 128 ms
6 iplsnng-chinng.abilene.ucaid.edu (198.32.8.77) [AS 11537] 144 ms 132 ms 136 ms
7 kscynng-iplsnng.abilene.ucaid.edu (198.32.8.81) [AS 11537] 152 ms 160 ms 140 ms
8 dnvrng-kscynng.abilene.ucaid.edu (198.32.8.13) [AS 11537] 164 ms 156 ms 152 ms
9 snvang-dnvrng.abilene.ucaid.edu (198.32.8.1) [AS 11537] 184 ms 176 ms 176 ms
10 pos-1-0.core0.eug.oregon-gigapop.net (198.32.163.17) [AS 4600] 192 ms 188 ms
  192 ms
11 uo-0.eug.oregon-gigapop.net (198.32.163.147) [AS 4600] 192 ms 200 ms 212 ms
12 ge-5-1.uonet1-gw.uoregon.edu (128.223.2.1) [AS 3582] 192 ms 188 ms
  ge-5-1.uonet2-gw.uoregon.edu (128.223.2.2) [AS 3582] 192 ms
13 darkwing.uoregon.edu (128.223.142.13) [AS 3582] 192 ms 188 ms 192 ms
```

- Now the path we see is **AS20965** (Geant), to **AS11537** (I2) to **AS4600** (the Oregon Gigapop) to **AS3582** (UO). If we checked other sites, we'd see still other paths, but in each case we could use the ASNs we see to compactly represent the path.

What Is An ASN?

- An Autonomous System Number is a number assigned to a group of network addresses managed by a particular network operator which share a common routing policy.
- Most ISPs, large corporations, and university networks have an ASN. For example, Google uses AS15169, Sprint uses AS1239, Intel uses AS4983, and so on. Some large networks with particularly complex routing policies may have multiple ASNs; others, with simple routing policies and only a single upstream network provider, may have none (their network blocks get announced using their upstream provider's ASN).
- You may want to think of an ASN as a number that "maps to" or represents a particular provider or network. ASNs are nice to work with because in most cases a given entity will only have one, no matter how many IP addresses or netblocks or customers they may have.

ASNs are New to Me. How Do I Translate the ASNs I See to Names?

- You can look ASNs up in the ARIN, RIPE, APNIC, LACNIC, AFRINIC, JPNIC, TWNIC (etc.) whois databases, just like IP addresses, either checking with a whois client or via the web whois interface provided by each of those registrars.
- If you don't find an ASN in the ARIN whois (for example), you may be redirected appropriately, or you may just need to try the other regions (e.g., check RIPE, check APNIC, check LACNIC, etc., etc.), until you finally get a match.
- Usually you'll preface the actual number with AS when looking it up, e.g., AS3582, but if you have difficulty getting a match with the AS included as a literal part of the query, try querying on just the actual AS number itself (this can help when the ASN you're trying to map is part of a range of ASNs documented via a single entry in the database).

Example of Looking Up an ASN

- Assume, for example, we want to know who owns AS20965:

```
% whois -h whois.ripe.net AS20965
[snip]
aut-num:        AS20965
as-name:        GEANT
descr:          The GEANT IP Service
[snip]
role:           DANTE Operations
address:        City House, 126-130 Hills Road
address:        Cambridge CB2 1PQ, UK
phone:          +44 1223 371300
fax-no:         +44 1223 371371
[snip]
```

The Origin AS; Detecting Hijacking

- Coming back to the traceroutes we did from Seattle and Switzerland, in each case the **last AS** in the path was the same: **AS3582**. That's the "origin AS."
- In our case, 128.223.142.13 belonged to UO and AS3582 also belonged to UO, so we can feel fairly comfortable that the 128.223.142.13 address was being used by an appropriate party. If bad traffic was seen from 128.223.142.13, UO should indeed be the ones to hear about it.
- But what if we'd seen some other AS other than 3582?
- **If/when a network address block gets hijacked, the ASN we'd normally expect to see ends up getting replaced with a different ASN, the ASN of the network that's injecting an unauthorized route for the hijacked netblock.**

Minimizing Your Address Hijacking Risk

- Ensure that whois information for your domain(s), your netblock(s), and (if applicable) your ASN is accurate and up to date.
- Be sure you or your networking engineering staff monitors your network blocks for potential hijacking. Some route monitoring and reporting projects include:
 - RIPE's myASN (<http://www.ris.ripe.net/myasn.html>)
 - UNM Internet Alert Registry
(<http://cs.unm.edu/~karlinjf/IAR/index.php>) and
 - Colorado State's Prefix Hijack Alert System
(<http://netsec.cs.colostate.edu/phas/>)

What If We Find That Our Address Space Does End Up Getting Hijacked?

- If you discover your address space is getting used without your permission, some steps you may want to take include:
 - contact the provider that's routing your block and ask them to stop (the problem may be a simple typographical error and they may be happy to correct the problem)
 - if that doesn't resolve the problem, contact the provider who's upstream of the ISP that's routing your address space, and ask them to intervene
 - a last resort can involve announcing so-called more specific routes (since traffic always follows the most specific route that's being announced, that can have the practical effect of pulling traffic back where it belongs, but there are a variety of pragmatic and philosophical reasons why you should avoid doing this if at all possible)

5. Sprites, Hobgoblins and Other Things We Didn't Cover Today

Wow, Are Those The Only Security Monsters Out There These Days?

- No. As we finish things up today, I don't want you to think that the four areas I highlighted are the only areas where serious security problems exist, because that's not the case -- four monsters were just all I figured I could fit in before I ran out of time. If it turns out we do still have some time and you all aren't too fatigued to continue, we could also talk about: (a) security policies (including regulatory compliance for things like CALEA); (b) personnel and the so-called insider threat; (c) physical security of your data center; (d) business continuity/disaster recovery planning; (e) wireless security; (f) password-related vulnerabilities; (g) spam; (h) phishing; (i) non-enterprise network security (e.g., things like the security of physical plant and alarm system networks); etc., etc., etc.
- What's uppermost on YOUR mind at this point?

Thanks for the Chance to Talk!

Are there any questions?