Security Now! #576 - 09-06-16

Flip Feng Shui

This week on Security Now!

The continuing woes of WoSign, autonomous micro-recon drones turn out to be real, a new crypto attack on short block ciphers prompts immediate changes in OpenVPN and OpenSSL, Introducing a new Security Now! Abbreviation: "YAWTTY" Yet Another Way To Track You, a discouraging social engineering experiment, another clever USB attack, a bunch of fun miscellany... and a look at the weaponizing of RowHammer with "Flip Feng Shui"... the most incredibly righteous and sublime hack.. ever! (And our follow-up to last week's Security Now! Puzzler.)

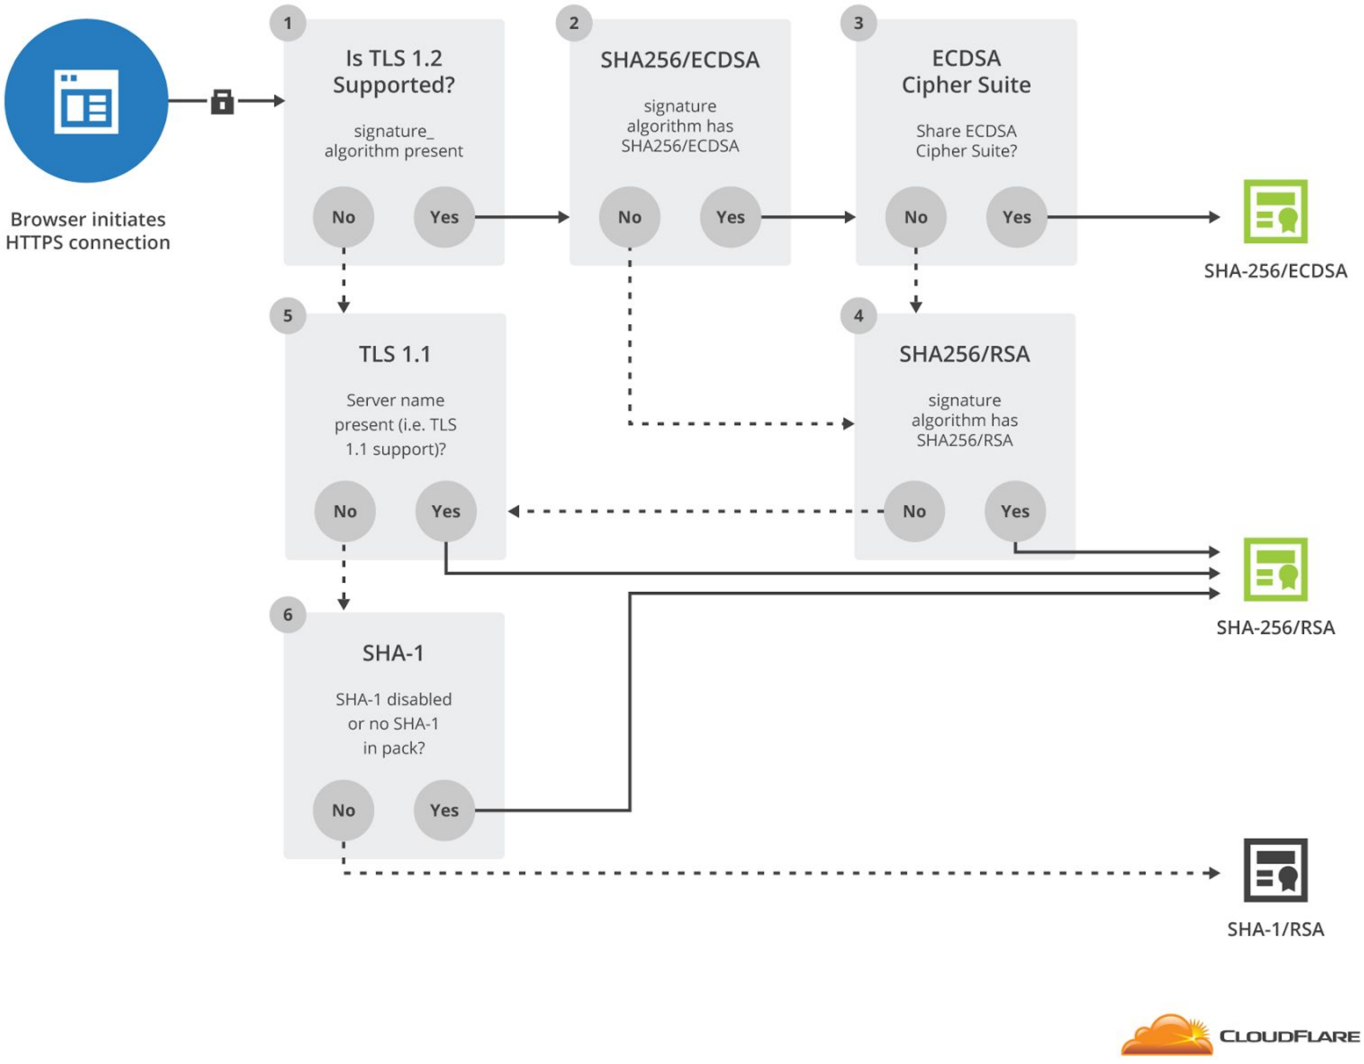
Browser initiates HTTPS connection
1
Is TLS 1.2 Supported?
signature_ algorithm present
No
Yes
2
SHA256/ECDSA
signature algorithm has SHA256/ECDSA
No
Yes
3
ECDSA Cipher Suite
Share ECDSA Cipher Suite?
No
Yes
SHA-256/ECDSA
5
TLS 1.1
Server name present (i.e. TLS 1.1 support)?
No
Yes
4
SHA256/RSA
signature algorithm has SHA256/RSA
No
Yes
SHA-256/RSA
6
SHA-1
SHA-1 disabled or no SHA-1 in pack?
No
Yes
SHA-1/RSA
CLOUDFLARE

Security News

Mozilla: Incidents involving the CA WoSign
https://groups.google.com/forum/#!msg/mozilla.dev.security.policy/k9PBmyLCi8I/mKSMaz9eCgAJ

Dear "Mozilla Dev Security Policy"...

Several incidents have come to our attention involving the CA "WoSign". Mozilla is considering what action it should take in response to these incidents. This email sets out our understanding of the situation.

Before we begin, we note that Section 1 of the Mozilla CA Certificate Enforcement Policy[0] says: "When a serious security concern is noticed, such as a major root compromise, it should be treated as a security-sensitive bug, and the Mozilla Policy for Handling Security Bugs should be followed." It is clear to us, and appears to be clear to other CAs based on their actions, that misissuances where domain control checks have failed, fall into the category of "serious security concern".

- First Incident (0):

  On or around April 23rd, 2015, WoSign's certificate issuance system for their free certificates allowed the applicant to choose any port for validation. Once validation had been completed, WoSign would issue certificates for that domain. A researcher was able to obtain a certificate for a university by opening a high-numbered port (>50,000) and getting WoSign to use that port for validation of control.

  This problem was reported to Google, and thence to WoSign and resolved. Mozilla only became aware of it recently.

  - Before the recent passage of Ballot 169 in the CAB Forum, which limits the ports and paths which can be used, the Baseline Requirements said that one acceptable method of domain validation was "Having the Applicant demonstrate practical control over the FQDN by making an agreed upon change to information found on an online Web page identified by a uniform resource identifier containing the FQDN". This method therefore did not violate the letter of the BRs. However, Mozilla considers the basic security knowledge that ports over 1024 are unprivileged should have led all CAs not to accept validations of domain control on such ports, even when not documented in the BRs.

  - The misissuance incident was not reported to Mozilla by WoSign as it should have been (see above).

  - This misissuance incident did not turn up on WoSign's subsequent BR audit[1].

- Second Incident (1):

  In June 2015, an applicant found a problem with WoSign's free certificate service, which allowed them to get a certificate for the base domain if they were able to prove control of a subdomain.

  The reporter proved the problem in two ways. They accidentally discovered it when trying to get a certificate for med.ucf.edu and mistakenly also applied for www.ucf.edu, which was approved. They then confirmed the problem by using their control of theiraccount.github.com/theiraccount.github.io to get a cert for github.com, github.io, and www.github.io.

  They reported this to WoSign, giving only the Github certificate as an example. That cert was revoked and the vulnerability was fixed. However recently, they got in touch with Google to note that the ucf.edu cert still had not been revoked almost a year later.

  - The lack of revocation of the ucf.edu certificate (still unrevoked at time of writing, although it may have been by time of posting) strongly suggests that WoSign either did not or could not search their issuance databases for other occurrences of the same problem. Mozilla considers such a search a basic part of the response to disclosure of a vulnerability which causes misissuance, and expects CAs to keep records detailed enough to make it possible.

  - This misissuance incident was not reported to Mozilla by WoSign as it should have been (see above).

  - This misissuance incident did not turn up on WoSign's subsequent BR audit[1].

- Third Incident (2):

  In July 2016, it became clear that there was some problems with the StartEncrypt automatic issuance service recently deployed by the CA StartCom. As well as other problems it had, which are outside the scope of this discussion, changing a simple API parameter in the POST request on the submission page changed the root certificate to which the resulting certificate chained up. The value "2" made a certificate signed by "StartCom Class 1 DV Server CA", "1" selected "WoSign CA Free SSL Certificate G2" and "0" selected "CA ", another root certificate owned by WoSign and trusted by Firefox.

  Using the value "1" led to a certificate which had a notBefore date (usage start date) of 20th December 2015, and which was signed using the SHA-1 checksum algorithm.

  - The issuance of certificates using SHA-1 has been banned by the Baseline Requirements since January 1st, 2016. Browsers, including Firefox, planned to enforce this[2] by not trusting certs with a notBefore date after that date, but in the case of Firefox the fix had to be backed out due to web compatibility issues. However, we are considering how/when to reintroduce it, and CAs presumably know this.

- The issuance of backdated certificates is not forbidden, but is listed in Mozilla's list of Problematic Practices[3]. It says "Minor tweaking for technical compatibility reasons is accepted, but backdating certificates in order to avoid some deadline or code-enforced restriction is not."

- WoSign deny that their code backdated the certificates in order to avoid browser-based restrictions - they say "this date is the day we stop to use this code"[4]. If that is true, it is not clear to us how StartCom came to deploy WoSign code that WoSign itself had abandoned.

- It seems clear from publicly available information that StartCom's issuance systems are linked to WoSign's issuance systems in some way. Nevertheless, it should not have been possible for an application for a cert from StartCom to produce a cert signed by WoSign.

- This misissuance incident was not reported to Mozilla by WoSign as it should have been.

Taking into account all these incidents and the actions of this CA, Mozilla is considering what action to take. Your input is welcomed.

- 155 posts from 34 authors

Micro-recon: US developing autonomous drones to sweep urban war zones
https://www.rt.com/viral/358206-military-tech-battlefield-drone/

Shield AI, a San Diego, California-based tech firm specializing in mini-reconnaissance quadcopters has been awarded a million dollar contract to provide drones that scour urban battlefields and beam back critical information to the US army. (Hopefully not U.S. domestic "urban battlefields".)

Shield AI’s mission statement suggests their flying machines can help solve the “intelligence deficit” that can often mean the difference between life and death for military personnel dropped into densely populated city war zones.

A notice on the US business procurement website, FedBizOpps, reveals the company has recently been contracted by the US Army and Naval Special Warfare Command to work on autonomous tactical airborne drones.

There is little detail in the $1 million contract awarded by the Defense Innovation Unit Experimental (DIUx) - the new tech-focussed outfit tasked with gaining the technological jump on America’s enemies - for a nine-month “prototype project in the area of Autonomous Tactical Airborne Drones.

A video sneak peek of Shield AI creations posted on the company website give some indication of their drones’ capabilities. The footage shows a micro drone being launched by tapping the screen of a smartphone, after which it takes off and, "with no remote control, no pilot, and full autonomy" maps narrow corridors of a building (looks like a storage rental facility) without human assistance.

In an interview, Shield AI's CEO, Ryan Tseng, explained that their mini-copters can be flown into an urban area and then used to “plan paths” through that environment. QUOTE: “Our mission is to protect our service members and innocent civilians with artificially intelligent systems. Basically we are packing all the capabilities of a Google self-driving car into a micro UAV or robot that fits in the palm of your hand. So it flies around all by itself, exploring... and its mission is to see everything without an operator actually having to control it.”

HTTPS and OpenVPN face new attack that can decrypt secret cookies
More than 600 sites found to be vulnerable to demanding exploit called Sweet32.
http://arstechnica.com/security/2016/08/new-attack-can-pluck-secrets-from-1-of-https-traffic-affects-top-sites/
Paper to be presented toward the end of October (23rd), next month, at the ACM Conference on Computer and Communication:

- FROM THE PAPER: Cryptographic protocols like TLS, SSH, IPsec, and OpenVPN commonly use block cipher algorithms, such as AES, Triple-DES, and Blowfish, to encrypt data between clients and servers. To use such algorithms, the data is broken into fixed-length chunks, called blocks, and each block is encrypted separately according to a mode of operation. Older block ciphers, such as Triple-DES and Blowfish use a block size of 64 bits, whereas AES uses a block size of 128 bits.

  It is well-known in the cryptographic community that a short block size makes a block cipher vulnerable to birthday attacks, even if there are no cryptographic attacks against the block cipher itself. We observe that such attacks have now become practical for the common usage of 64-bit block ciphers in popular protocols like TLS and OpenVPN. Still, such ciphers are widely enabled on the Internet. Blowfish is currently the default cipher in OpenVPN, and Triple-DES is supported by nearly all HTTPS web servers, and currently used for roughly 1-2% of HTTPS connections between mainstream browsers and web servers.

  We show that a network attacker who can monitor a long-lived Triple-DES HTTPS connection between a web browser and a website can recover secure HTTP cookies by capturing around 785 GB of traffic. In our proof-of-concept demo, this attack currently takes less than two days, using malicious Javascript to generate traffic. Keeping a web connection alive for two days may not seem very practical, but it worked easily in the lab. In terms of computational complexity, this attack is comparable to the recent attacks on RC4. We also demonstrate a similar attack on VPNs that use 64-bit ciphers, such as OpenVPN, where long-lived Blowfish connections are the norm.

  Countermeasures are currently being implemented by browser vendors, OpenSSL, and the OpenVPN team, and we advise users to update to the latest available versions.

So... The attack requires both that:

- An attacker has the ability to monitor traffic passing between the end user and a vulnerable website.

- And also to have injected a JavaScript exploit engine onto a webpage loaded by the user's browser. This must be done either by actively manipulating an HTTP response on the wire or by hosting a malicious website that the user is tricked into visiting.

The JavaScript then spends the next 38 hours collecting about 785GB worth of data to decrypt the cookie, which allows the attacker to log into the visitor's account from another browser.

Attacking OpenVPN requires 18 hours and 705GB of data to recover a 16-byte authentication token.

Time to retire the short bit block ciphers: Triple DES in SSL and OpenVPN's use of Blowfish. Modern 128-bit block ciphers such as the AES standard Rijndael are fine.

In response to this attack:

- OpenVPN just released an update which actively discourages the use of 64-bit ciphers.

- v2.3.12 "This release includes many small improvements and fixes. This is the first release that actively discourages the use of 64-bit block ciphers for security reasons."

- OpenSSL maintainers plan to disable Triple DES in v1.1.0 and to deprecate the claimed security level of 3DES from "high" to "medium" to bias the auto-selection logic against choosing it.

This is security research working correctly. There will always be similar discoveries used against those trusting the technology. But this is why we MUST keep the research process available and open.

How Your Smartphone Light Sensor Could Help Websites Track You
YAWTTY - Yet Another Way To Track You.
https://motherboard.vice.com/read/smartphone-light-sensor-tacking-privacy
https://blog.lukaszolejnik.com/privacy-of-ambient-light-sensors/
https://www.w3.org/TR/ambient-light/

- This specification defines a concrete sensor interface to monitor the ambient light level or illuminance of the device’s environment.

- The Ambient Light Sensor extends the Generic Sensor API [GENERIC-SENSOR] to provide information about ambient light levels, as detected by the device’s main light detector, in terms of lux units. The light-level media feature [MEDIAQUERIES-4] provides less granular information about the ambient light level.

- Can also determine the exact COLOR of the light (individual R/G/B Lux levels.)

- Yet another way to bleed a bit of identifying information from a user.

So much for counter-phishing training: Half of people click anything sent to them
(Even people who claimed to be aware of risks clicked out of curiosity.)
http://arstechnica.com/security/2016/08/researchers-demonstrate-half-of-people-will-click-on-any-link-theyre-sent/

- Social Engineering is alive and well…
- 1700 University students who claimed to be aware of the risks of unknown links.
- The e-mail and Facebook accounts were set up with the ten most common names in the age group of the targets.
- The Facebook profiles had varying levels of publicly accessible profile and timeline data—some with public photos and profile photos, and others with minimal data.
- The messages claimed the links were to photos taken at a New Year's Eve party held a week before the study. (So... high relevance)
- Two sets of messages were sent out: in the first, the targets were addressed by their first name; in the second, they were not addressed by name, but more general information about the event allegedly photographed was given. Links sent resolved to a webpage with the message "access denied," but the site logged the clicks by each student.
- The messages that addressed the targets by name scored clicks from 56 percent of e-mail targets and 37 percent of Facebook message recipients.
- The less-targeted messages in the second test only yielded 20 percent results for the e-mails, but they scored 42 percent via Facebook messages.
- The German security researchers who conducted the study said: "The overall results surprised us, as 78 percent of participants stated in the questionnaire that they were aware of the risks of unknown links, and only 20 percent from the first study and 16 percent from the second study said that they had clicked on the link."
- But of those claiming they were security savvy they found that 45 and 25 percent respectively had clicked on the links.

There's a new PC / Mac hack on the scene…

- Once again, utilizes USB's Plug & Play mechanism to install a new network adapter. If this adapter provides DCHP services (to provide WPAD (proxy auto-config) and DNS)
- “DHCP has a higher priority than DNS: if DHCP provides the WPAD URL, no DNS lookup is performed.” (from the Wikipedia article on WPAD)
- Responder
    - https://github.com/SpiderLabs/Responder
    - Responder is a LLMNR, NBT-NS and MDNS poisoner, with built-in HTTP/SMB/MSSQL/FTP/LDAP rogue authentication server supporting NTLMv1/NTLMv2/LMv2, Extended Security NTLMSSP and Basic HTTP authentication.

Miscellany

Spaces or Tabs??
A Googler analyzed a billion files to settle the programming dispute made famous by HBO's 'Silicon Valley'
https://amp.businessinsider.com/google-silicon-valley-tabs-spaces-debate-2016-8

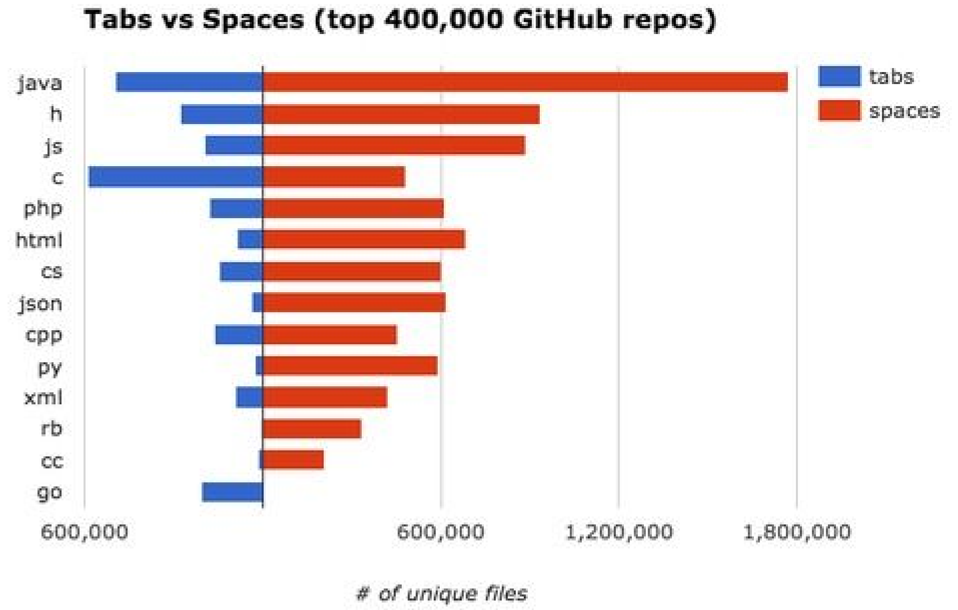

As we see... Spaces dominate except for C and Go.

John Carmack

- Wikipedia: "John D. Carmack is an American game programmer, aerospace and virtual reality engineer. He co-founded ID Software. Carmack was the lead programmer of the id video games Commander Keen, Wolfenstein 3D, Doom, Quake, Rage and their sequels. In August 2013, Carmack took the position of CTO at Oculus VR
- @ID_AA_Carmack
  Create and run an empty activity project in Android Studio, and I get a 38 MB folder with 1,175 files. We've just given up on elegance.

Andrew Hutcheson @AndrewHutchATX

- @SGgrc @leolaporte "I'm not scared of a computer passing the turing test... I'm terrified of one that intentionally fails it."

Adam Stearn @MonkeyThink

- @SGgrc Steve - what was that mail archiving application you mentioned weeks ago on SN?
- It's "Mailstore Home". I'm using v8, but they're at v9 now. Great tool. Using it continually. And it's FREE for personal use.

ZendoDeb @ZendoDeb

- @SGgrc #3: given that Apple refuses to allow outside security audits, what is the basis for saying iOS is "most secure mobile platform?"
- A: As we know, Apple exercises extreme control over their closed iDevice ecosystem. Google currently has much less control. This gives Android users more freedom... but at a cost in device security. Though Google is much better with their own devices, the vast majority of Android Smartphones worldwide are either never patched or are patched partially and/or late.

Tim Stewart @TimAStewart

- @GibsonResearch is loading a form over HTTP and posting to HTTPS a security risk??
- Yes... because a form over HTTP could have been altered in ANY way if it was delivered without security. So, for example, it could "double-send" the filled-in data, both to its original destination and to another (by sending it in a URL link query, for example.)

SpinRite

- sjh_canada @sjh_canada
  @SGgrc @leolaporte
  EXE is short for "EXEcution time binding"
  COM is short for "COMpile time binding" - data & vars in a 64k memory section

---

Flip Feng Shui

The most incredibly righteous & sublime hack!!
New cloud attack takes full control of virtual machines with little effort
(Existing crypto software "wholly unequipped" to counter Rowhammer attacks.)
http://arstechnica.com/security/2016/08/new-attack-steals-private-crypto-keys-by-corrupting-data-in-computer-memory/

Links:

- https://www.vu.nl/en/news-agenda/news/2016/jul-sep/new-hacking-technique-imperceptibly-changes-memory-virtual-servers.aspx
- https://www.ncsc.nl/english/current-topics/news/researchteam-presents-flip-feng-shui-attack-method-at-usenix-security-symposium-2016.html
- https://www.vusec.net/projects/flip-feng-shui/
- http://www.cs.vu.nl/~herbertb/download/papers/flip-feng-shui_sec16.pdf

A review of DRAM and the RowHammer attack

- Memory Disturbance Errors (March 10, 2015)
- Double Row Hammer
- Mitigations in DDR3 and DDR4 DRAM

A review of Mapped Memory Management

- The "deck of cards" analogy
- The view of memory that an application “sees” bears ZERO resemblance to actual physical memory. (At the block (page) level.)

KSM: Kernel Same-Page Merging

- To optimize a large system's memory, the Virtual Machine Manager periodically scans memory searching for identical pages of memory and, when found:
  - Marks them "read-only" to detect any changes, then
  - MERGES THEM by updating all VM's having an exact copy of that page to point to the SAME single copy of the main memory.

Putting it all together:

- Survey the application’s entire memory allocation for qualifying flippable bits.
  - “Qualifying” means that the bit occurs at a useful offset in the page.
- Create a duplicate of the target page in the victim VM.
- Flip the bit in our mirrored page.
-

Flipping a single bit of an RSA key renders the result VASTLY more factorable!!

Last week's puzzler:

- If possible, spoof the date.
  - If the device is only using the local clock:
    - Set the clock back... but not too far: In addition to a "not valid after", certificates also have a "not valid after."

- If the device obtains time from the network:
  - Look at the device's traffic to see which NTP servers it queries.
  - Aim the device's DNS to a server which spoofs the real NTP to a "retrograde" NTP server.

- Deploy Cloudflare's "No Browser Left Behind" solution to serve the public domain where the devices connect.
  - NBLB which will detect that the SSL/TLS handshake is pre TLS v1.1 and will dynamically serve a valid SHA-1 certificate.
  - https://blog.cloudflare.com/tls-certificate-optimization-technical-details/