Security Now! #578 - 09-20-16

GRC's XSS Adventure

This week on Security Now!

Concerns over a significant expansion in effectively warrantless intrusion into end-user computers, the forthcoming change in Internet governance, and generation of a shiny new (and bigger) DNSSEC root signing key, Google's next move in using Chrome to push for improved security, the interesting details emerging from a successful NAND memory cloning attack on the iPhone 5c, some fun miscellany... and Steve shares the details and findings of a recent Cross-Site Scripting (XSS) problem on GRC, including the best website security scanner he found and now recommends!

Security News

Amendments to Rule 41 of the Federal Rules of Criminal Procedure

- The proposal, which will automatically go into effect on December 1st unless Congress passed legislation to block this proposal, would allow the government to hack any number of computers -- mllions or more -- with a single warrant. With just six work weeks remaining on the Senate schedule and a long Congressional to-do list, time is running out.

- The government says it needs this power to investigate "botnets" -- as we know so well, networks of devices infected with malware and controlled by a criminal. But the Justice Department has given the public very little information about its hacking tools, and how it plans to use them. And the amendments to Rule 41 are woefully short on protections for the security of hospitals, life-saving computer systems, or the phones and electronic devices of innocent bystander Americans.

- No one believes the government is setting out to damage victims' computers. But history shows just how hard it is to get hacking tools right. Indeed, recent experience shows that tools developed by law enforcement have actually been co-opted and used by criminals and miscreants.

- Why would Congress approve such a short-sighted proposal? It didn't. Congress had no role in writing or approving these changes, which were developed by the US court system through an obscure procedural process. This process was intended for updating minor procedural rules, not for making major policy decisions.

- Wired Magazine editorializes, stating: This kind of vast expansion of government mass hacking and surveillance is clearly a policy decision. This is a job for Congress, not a little-known court process. If Congress had to pass a bill to enact these changes, it almost surely would not pass as written. The Justice Department may need new authorities to identify and search anonymous computers linked to digital crimes. But this package of changes is far too broad, with far too little oversight or protections against collateral damage.

- What exactly is being changed?

    - The first would let magistrate judges issue search warrants to remotely search—essentially hack—computers outside their jurisdiction if the location of the computer has been intentionally concealed through technical means. Currently magistrates can only issue warrants to search and seize property within their court's jurisdiction. The pending change would mean that when a hacker or child pornographer uses Tor or some other proxy to conceal their real IP address and location, law enforcement would not be required to determine the location of the computer to get permission to hack it.

    - [Steve's comment:] And any member of a botnet, using spoofed source IPs, is also intentionally hiding its identity and location.

    - The second amendment would let magistrates issue a warrant outside their jurisdiction when the computers to be searched are part of a cybercrime investigation—as defined by the Computer Fraud and Abuse Act—have been "damaged without authorization" and "are located in five or more districts." The rules committee says the amendment is intended to "eliminate the burden of attempting to secure multiple warrants in numerous districts" and allow a single judge to oversee an investigation. But the description of the computers to be searched has nothing to do with criminal suspects, critics point out, and instead refers to victims' computers.

    - The third Rule 41 change is even more tricky. Law enforcement has to find a way to tell people when a search of their property has occurred. With in-person searches, this is easy to do. They either hand notice "to the person from whom, or from whose premises, the property was taken" or leave a notice "at the place where the officer took the property." But this is challenging with remote searches when the computer's "place" and computer owner are unknown. Under the amendment, law enforcement "must make reasonable efforts" to serve a copy of the warrant on the person whose property was searched, which "may be accomplished by any means, including electronic means."

    - This concerns civil liberties groups, since an email notification or pop-up message from law enforcement could easily look like a phishing attack to a botnet victim and be ignored. Enterprising hackers would also adopt this as a tactic to trick users into clicking on malicious links or attachments.

- Links:
  https://www.wired.com/2016/05/now-government-wants-hack-cybercrime-victims/
  https://www.wired.com/2016/09/government-will-soon-able-legally-hack-anyone/

NTIA's contract with ICANN to handle IANA is expiring in ten days!

- IANA - Internet Assigned Numbers Authority, which includes DNS.
  - While the Internet is unique in being a worldwide network that's largely free from central coordination, there is a technical need for the definition and management of communications standards which requires network wide, thus global, coordination. This coordination role is provided by the Internet Assigned Numbers Authority.
    - IANA manages the DNS Root, the .int and .arpa domains.
    - It coordinates the global pool of IP and AS (autonomous system) numbers.
    - It maintains the Internet Protocol Assignments.

- ICANN -- the Internet Corporation for Assigned Names and Numbers -- is a multinational multistakeholder body, based in Los Angeles, with many member countries including China and Russia. While the Internet has been growing and "happening", ICANN has managed the IANA functions under a contract with the Commerce Department's NTIA - National Telecommunications and Information Administration. However, the United States has long made clear that it intended to privatize the DNS in order to facilitate international participation in its management.

- Now the NTIA intends to allow its contract with ICANN to expire on Sept. 30, at which time ICANN will assume stewardship of the IANA's key technical functions.

  Most internet experts support the internet governance transition, because it counters the growing argument repressive regimes use to lobby for greater power over internet governance, or break off from the global internet altogether. Everyone loses if the Internet becomes fragmented.

- So, the IANA will be moving from previous contractual control by the U.S. control to ICANN, the international body.

Meanwhile... the DNS Root is being re-keyed/re-signed...

- The DNSSEC key pair at the top of the DNS -- the Root Zone Signing Key -- is being changed for the first time.
- There are no known problems, it's just good security policy.
- The signing key's length is being doubled from 1024 to 2048 bits.
- The process:
  - ICANN incorporates some extraordinary security measures, and considers its potential threats as everything up to nation states. At its quarterly ceremonies, so-called "crypto officers" from all over the world congregate in one of the key management facilities, after passing layers of physical and digital security.

- This October, in a hyper-secure key management facility on the US east coast, ICANN will generate a new cryptographic key pair. One half of that pair is private, and will be kept -- super securely -- by ICANN; the other is public. Internet service providers, hardware manufacturers, Linux and other OS developers -- anything and anyone who needs to verify future DNSSEC records, needs to have the updated DNSSEC public key.

- In the first quarter of next year (2017), two employees will transport a copy of the encrypted key files on a smartcard over to another facility on the west coast, using regular commercial transport. Eventually, the public part of the key pair will be distributed to other organizations.

- In all, the whole switchover will take around two years from start to finish. The new key will appear in the DNS for first time on July 11, 2017. In October 2017, the new key will be used for making signatures.

Google uses Chrome to further enforce browser security

- https://security.googleblog.com/2016/09/moving-towards-more-secure-web.html
- Starting in January (2017), Google's Chrome web browser, will treat any website encoded in HTTP as a non-secure site if it transmits credit card information or passwords.
- <quote> To help users browse the web safely, Chrome indicates connection security with an icon in the address bar. Historically, Chrome has not explicitly labelled HTTP connections as non-secure. Beginning in January 2017 (Chrome 56), we'll mark HTTP pages that collect passwords or credit cards as non-secure, as part of a long-term plan to mark all HTTP sites as non-secure.

  Chrome currently indicates HTTP connections with a neutral indicator. This doesn't reflect the true lack of security for HTTP connections. When you load a website over HTTP, someone else on the network can look at or modify the site before it gets to you.

  A substantial portion of web traffic has transitioned to HTTPS so far, and HTTPS usage is consistently increasing. We recently hit a milestone with more than half of Chrome desktop page loads now served over HTTPS. In addition, since the time we released our HTTPS report in February, 12 more of the top 100 websites have changed their serving default from HTTP to HTTPS.

  Studies show that users do not perceive the lack of a "secure" icon as a warning, but also that users become blind to warnings that occur too frequently. Our plan to label HTTP sites more clearly and accurately as non-secure will take place in gradual steps, based on increasingly stringent criteria. Starting January 2017, Chrome 56 will label HTTP pages with password or credit card form fields as "not secure," given their particularly sensitive nature.

  In following releases, we will continue to extend HTTP warnings, for example, by labelling HTTP pages as "not secure" in Incognito mode, where users may have higher expectations of privacy. Eventually, we plan to label all HTTP pages as non-secure, and change the HTTP security indicator to the red triangle that we use for broken HTTPS.

The bumpy road towards iPhone 5c NAND mirroring

- http://arxiv.org/abs/1609.04327
- http://arxiv.org/pdf/1609.04327v1

ABSTRACT - This paper is a short summary of a real world mirroring attack on the Apple iPhone 5c passcode retry counter under iOS 9. This was achieved by desoldering the NAND Flash chip of a sample phone in order to physically access its connection to the SoC and partially reverse engineering its proprietary bus protocol. The process does not require any expensive and sophisticated equipment. All needed parts are low cost and were obtained from local electronics distributors. By using the described and successful hardware mirroring process it was possible to bypass the limit on passcode retry attempts. This is the first public demonstration of the working prototype and the real hardware mirroring process for iPhone 5c. Although the process can be improved, it is still a successful proof-of-concept project. Knowledge of the possibility of mirroring will definitely help in designing systems with better protection. Also some reliability issues related to the NAND memory allocation in iPhone 5c are revealed. Some future research directions are outlined in this paper and several possible countermeasures are suggested. We show that claims that iPhone 5c NAND mirroring was infeasible were ill-advised.

<quote cool tech details> The process of cloning involves creating a fully working copy of the NAND Flash memory chip. However, as it was already mentioned in the previous section, simply copying the 8GB information from the original chip into another identical chip taken from other iPhone 5c does not give the desired result and the iPhone does not boot.

Some additional research was undertaken to figure out why simple copying does not work. For that the same model of the NAND Flash chip was programmed with the data from the original chip and then the communication was analyzed with both an oscilloscope and a logic analyzer. First, some pages were accessed from addresses outside the normal 16GB space. For example, instead of reading and writing to the block 0x00041Axx the CPU was accessing the block 0x00841Axx. Although such accesses [wrap around and] are mapped back into the 0x00041Axx block, the page numbers were different as well as the status of those pages. Figure 20 shows status and checksums for the first several pages for the blocks 0x00041Axx and 0x00841Axx. It can be noted that the page 0x00841A02 is mapped to 0x00041A03, page 0x00841A03 to 0x00041A05 and so on. The status of the pages for 0x00041Axx block was 0x61 in comparison with 0x40 for 0x00841Axx block. Second, some irregularities were found in the communication prior the access to those hidden pages. Although the data transfer during the access is performed in the SDR mode at 17MHz, while the configuration commands use even the slower speed of about 1MHz, some data inside the commands are smuggled at an astonishing rate of 256MB/s in DDR3 mode. Also, a dummy value for data bit 7 was introduced for the period of 23ns. Given that the data setup time in those transfers is less than 1ns there is a very high chance that such information would be overlooked by most would-be attackers.

After those findings the implementation of the communication protocol was amended in the test board. Then the data-mirroring software was modified to include cloning the hidden pages. As a result the newly created clone of the original NAND chip was fully functional in the iPhone 5c. It was then tested with six incorrect passcode attempts before replacing it with the original chip. After the boot process it was possible to enter the incorrect passcodes again six times until the one minute delay was introduced. This fully proved the correctness of the hardware NAND mirroring attack on iPhone 5c.

Because there is no limitation on the number of such NAND clones, they can be created in advance and restored in parallel when one of them is being used for passcode testing. This way it only requires 45 seconds per six passcode attempts. For 4-digit passcode the maximum attack time would be (10000/6) × 45 = 75000 seconds or about 20 hours. For 6-digit passcode this time will increase to about 3 months which in some cases might be acceptable.

[Steve's note:] Note that the "vulnerability" as such was due to the presence of a "bus" -- i.e. separate communicating subsystems. If/when Apple further integrates the memory that opportunity to observe communications will disappear. Or Apple could explicitly encrypt that comm bus with a private key known to each end point.

Miscellany

- Will D. (@woden325) 9/20/16, 6:55 AM
    - @SGgrc My school's DNS Server was down this morning, but I remembered your DNS Benchmark tool. I'm back online & faster than before. Thanks!

- Kevin Garman (@kevinrgarman) 9/15/16, 7:13 PM
    - @SGgrc You mentioned DynDNS...I use http://freedns.afraid.org/ free service.
    - FreeDNS - Free subdomain AND domain hosting!
    - Get "yourname.afraid.com" and it can point to your home IP.
    - Why is it free? It's quite simple. We wanted a challenge... that's it.

- Simon Zerafa (Simon Zerafa) 9/20/16, 5:21 AM
    - <quote> Hi Steve, Can I please request a PSA for this weeks show? My wife (Bethan) had a serious medical emergency yesterday and we hadn't filled in the Emergency Contact information on the iOS Health app. As a result, the emergency services were delayed in contacting me. Can we ask listeners with iOS devices to fill that info in and enable it on their lock screens? It should make it easier in the future for emergency services to contact people and won't require them to enter their iOS PIN code to access their emergency contact details If Leo is up for it then perhaps he can repeat this on the other shows also.

SpinRite

Michael Nation in Michigan
:
Steverino: (apologies to Steve Allen)
I run SpinRite on all my drives all the time and nothing ever goes wrong! I'm disappointed. If I'm paying for health insurance, I darn well expect to get sick!

(Steve notes: The health insurance analogy would be more like using RAID to RECOVER from drive problems. Using SpinRite is more like taking adequate vitamin D to PREVENT any drive problems in the first place!)

Anonymous Sender
Location: Somewhere
Subject: Sprinrite Rescues a Stuck WX Update
Date: 13 Sep 2016 21:40:39
:
My Windows 10 Version 1607 patches got stuck today, and this isn't the first time in all, although it is the first time for a version 1607 update. Retrying didn't and doesn't help. So usually I would just revert using system restore and try re-updating again. But this time I tried SpinRite... and IT WORKED!

GRC's XSS Adventure

Open Bug Bounty
https://www.openbugbounty.org/
https://www.openbugbounty.org/incidents/180211/

Acunetix
- Offer both an online scan and a local Windows app.
- Authenticating domain ownership was fast and easy and just worked.
- The local app was not well written. Its UI locks up during scan due to a lack of an independent UI thread. (Neither the PAUSE nor STOP buttons work while the scans are running.) Not particularly confidence inspiring.
- Both scans warned of dire findings and that the sky was falling... but they wouldn't elaborate -- at all -- without first receiving payment in full. So it was impossible to "evaluate" their findings... if, indeed, they had any.
- Annoying tease and heinously expensive -- per year.
- After a few days I began getting spammed with four or five eMails/day.

Beyond Security
- Unable to authenticate without their help.
- Required FLASH for displaying summary graphs
- Very nice people, but their scanner didn't find any problems.
- Gave GRC a score of 100.00 and an A+ grade.

Tinfoil Security
- https://www.tinfoilsecurity.com
- As with Beyond Security, the web-based domain ownership failed, presumably because they don't follow HTML 301 Redirects. BUT!... they also offered a DNS TXT record authentication... which worked beautifully.
- 3 month trial -- 100% functional and comprehensive
- Reported 64 problems

What did Tinfoil find?

- 64 problems - Mostly info, some low-impact, some they thought were high impact.

- Credit card number disclosure
  "4207 0610 0000 0013" on the PDP-8 pages -- Octal machine code.

- 10 instances of "Missing Subresource Integrity Protection"
  All externally loaded resources must have their content pinned using the subresource integrity mechanisms provided by modern browsers. This involves computing a hash of the contents of the resource, and specifying this hash when loading that resource. In the case of a script, this might look like the following:

  ```
  <script src="https://example.com/include.js"
          integrity="sha256-Rj/9XDU7F6pNSX8yBddiCIIS+XKDTtdq0//No0MH0AE="
          crossorigin="anonymous"></script>
  ```

- 14 eMail addresses, mostly in Security Now! transcripts.

- 2 "Found Robots.txt"
  If you require it, try not to list important files or directories in robots.txt; instead, password protect them such that even if they are crawled the crawler gets nothing but an authentication page. A good robots.txt file includes content like image directories or locations that are generated dynamically and do not work if a search engine accesses the page, but does not include administrative areas or server logs.

- 7 "Found an HTML Object"
  This is nothing to be concerned about at the moment, as we are purely providing it for informational purposes. It often gives the hacker a beachhead of where to begin searching for a vulnerability, but isn't a vulnerability in and of itself.

- 4 "Insecure Cookies"
  Set the 'Secure' flag in the cookie, such that all cookies are served over a secure channel like HTTPS only. If you tell us what software stack you're running on your website, we'll be able to give you more detailed results on how to fix this issue. You can do this from your Dashboard.

- 4 "Non HTTP-Only Cookies"
  Set the 'HttpOnly' flag in the cookie, such that cookies cannot be manipulated via client-side code like JavaScript. If you tell us what software stack you're running on your website, we'll be able to give you more detailed results on how to fix this issue. You can do this from your dashboard.

- 16 "Private IP address disclosures"
  These were, again, mostly in Security Now podcast transcripts. Remove private IP addresses from the body of the HTML pages. These can often be leftover from testing and may indicate that your website is not running in the environment it expects. These IP addresses also give information to an attacker they can be used as a beachhead for other

more harmful actions.

BUT... they DID find three things:

On GRC's HTTPS Fingerprints page

- CSRF (Cross-Site Request Forgery) False positive due to no replay protection Parsing out the domain name automatically protected the form

GRC's Web News Interface

- Old CGI executable echoed back the "utag" it is given

Our Transaction ID lookup

- I was returning the erroneous TransID and it COULD execute JS.

Lookup Transaction ID:
1'"()><ScRiPt >alert('-XSS-')</ScRiPt>
1'%22()><ScRiPt%20>alert('XSS')</ScRiPt>

Unsafe HTML character entities
< <
> >
& &
" "
' '
/ /

CSP - Content Security Policy
default-src 'none'
form-action https://*.grc.com https://grc.com
img-src https://*.grc.com https://grc.com
style-src https://*.grc.com https://grc.com
upgrade-insecure-requests https://*.grc.com https://grc.com

IMPLIED:
child-src 'none'
connect-src 'none'
font-src 'none'
media-src 'none'
object-src 'none'
plugin-types 'none'
style-src 'none'
NOT IMPLIED:
base-uri
form-action
frame-ancestors 'none'
plugin-types
report-uri
sandbox