Security Now! #569 - 07-19-16

Messenger, CryptoDrop & Riffle

This week on Security Now!

- A bit of daylight on the password sharing question, the trouble with self-reporting security breaches, trouble in TOR-land, what future AI assistants mean for our privacy, a terrific looking new piece of security monitoring freeware, a startlingly worrisome 20 year old fundamental Windows architectural design flaw, a problem with Juniper router's OS certificate validation, some errata, a bunch of miscellany, and follow-up dissection of Facebook Messenger's extra features, the anti-ransomware CryptoDrop, and MIT's "Riffle" anonymity-enforcing networking solution.

All Your IoT devices are doomed

Security News

Ninth Circuit Panel Backs Away From Dangerous Password Sharing Decision—But Creates Even More Confusion About the CFAA
https://www.eff.org/deeplinks/2016/07/ninth-circuit-panel-backs-away-dangerous-password-sharing-decision-creates-even

EFF: Three judges of the Ninth Circuit Court of Appeals have taken a step back from criminalizing password sharing, limiting the dangerous rationale of a decision issued by a panel of three different judges of the same court last week. That's good, but the new decision leaves so many unanswered questions that it's clear we need en banc review of both cases—i.e., by 11 judges, not just three—so the court can issue a clear and limited interpretation of the notoriously vague federal hacking statute at the heart of both cases, the Computer Fraud and Abuse Act (CFAA).

To recap, the court's language in last week's case, U.S. v. Nosal, was so broad that it seemed to make it a federal crime to use someone else's password, even with their knowledge and permission. In the new decision, in a case called Facebook v. Power Ventures, a separate Ninth Circuit panel acknowledged that a computer user can provide another person with valid authorization to use their username and password. That's the good news. But the decision leaves unanswered so many other questions about how the law can be interpreted, and its rationale is so confusing, that it's an invitation for more dangerous litigation and prosecutions under the CFAA.

The CFAA makes it illegal to engage in "unauthorized access" to any computer connected to the Internet. But the statue doesn't say what "authorized access" means or make clear where authorization must come from.

FDIC was hacked by China, and CIO covered it up
http://arstechnica.com/security/2016/07/fdic-was-hacked-by-china-and-cio-covered-it-up/
Sean Gallagher reporting last Wednesday:
A report published by the House Committee on Science, Space and Technology today found that hackers purported to be from China had compromised computers at the Federal Deposit Insurance Corporation repeatedly between 2010 and 2013. Backdoor malware was installed on 12 workstations and 10 servers by attackers—including the workstations of the chairman, chief of staff, and general counsel of FDIC. But the incidents were never reported to the US Computer Emergency Response Team (US-CERT) or other authorities, and were only brought to light after an Inspector General investigation into another serious data breach at FDIC in October of 2015.

The FDIC failed at the time of the "advanced persistent threat" attacks to report the incidents. Then-Inspector General at FDIC, Jon Rymer, lambasted FDIC officials for failing to follow their own policies on breach reporting. Further investigation into those breaches led the committee to conclude that former FDIC CIO Russ Pittman misled auditors about the extent of those breaches, and told employees not to talk about the breaches by a foreign government so as not to ruin FDIC Chairman Martin Gruenberg's chances of confirmation.

- "Self reporting" is, of course, a problem. If your boss instructs you not to report, on threat of termination, what do you do? Sony and RSA both had APT -- advanced persistent threats -- that hurt them significantly.

One of TOR's earliest contributors, Lucky Green, exits the TOR project
Here's the notice he posted: https://trac.torproject.org/projects/tor/ticket/19690

Dear friends,

Given recent events, it is no longer appropriate for me to materially contribute to the Tor Project either financially, as I have so generously throughout the years, nor by providing computing resources. This decision does not come lightly; I probably ran one of the first five nodes in the system and my involvement with Tor predates it being called "Tor" by many years.

Nonetheless, I feel that I have no reasonable choice left within the bounds of ethics, but to announce the discontinuation of all Tor-related services hosted on every system under my control.

Most notably, this includes the Tor node "Tonga", the "Bridge Authority", which I recognize is rather pivotal to the network

Tonga will be permanently shut down and all associated crytographic keys destroyed on 2016-08-31. This should give the Tor developers ample time to stand up a substitute. I will terminate the chron job we set up so many years ago at that time that copies over the descriptors.

In addition to Tonga, I will shut down a number of fast Tor relays, but the directory authorities should detect that shutdown quickly and no separate notice is needed here.

I wish the Tor Project nothing but the best moving forward through those difficult times,

--Lucky

- The IP address of Lucky's "Bridge Authority" node is hard coded into TOR apps, so over time they will likely need to be updated to reflect the network's new topology.

- Questions asking for more details about Lucky's decision have been unanswered.

- Early last month, another major TOR project figure, Jacob Appelbaum, decided to leave the project under some controversy.

Why the top 5 tech companies are dead set on AI
https://techcrunch.com/2016/07/18/why-the-top-5-tech-companies-are-dead-set-on-ai/

- The pace of smartphone and desktop hardware innovation has slowed.
- And this is inevitable. The way the world works. Once we figured out that most cars should have four wheels we were done with that question.
- The GUI Mouse & Cursor model. / Word Processing.
- The frontier appears to be AI:
    - Apple: Siri
    - Google: Assistant
    - Amazon: Alexa
    - Microsoft: Cortana
    - Facebook: Chatbots

- But... for AI, CONTEXT is king…

- Early speech recognition software needed to be trained by having the user read a word list.

- Handwriting recognition needs to adapt to and learn the unique style of its user.

- Because those are "mechanical" things, we don't think in terms of privacy... but when you stop to think about it, it was early biometric parameter acquisition.

- Now we're moving into the "personal assistant" space and what's the context against which our needs will be understood? The context will be: Everything available about us. Our entire lives.

- Our lives and habits and interests are going to be recorded and profiled and modelled in order for the next generation of personal digital assistants to be able to perform their jobs.

- It's happening now.

- Perhaps we won't notice or care... the way we didn't notice or care when we were training speech and handwriting recognition, because the benefit seemed to clearly outweigh the cost. But that may not be true for all of us.

"CertWatch" by (Beau) Blaser Software
http://www.blaser.us/software/certwatch/

http://bit.ly/sn569

SN #551 - March 17th:

- Leo: It's essentially a man in the middle sitting on my own machine, much like the certs from antivirus companies. Do you know of any tools that monitor cert stores and report when changes are made?

- Steve: So, if I weren't so far behind, I would immediately whip out a utility. I'm not going to.

- Leo: That's a great idea, yeah.

- Steve: It **is** a great idea.

- Leo: Someone should write that.

- Steve: And that's why I put this in this Q&A. Remember that Mark Russinovich just recently updated the Sysinternals tool so that, with a command line, it'll do that. Somebody could write a little Python front end to the Sysinternals tool that is invoked by the scheduler, or maybe just runs in the background and checks every day. If anyone does, make sure you send me a note, and I will make you famous. I will tell everybody about it because that ought to exist.

Enter: CertWatch
http://www.blaser.us/software/certwatch/
http://bit.ly/sn569 (Google hasn't yet found it.)
(No hyphen this week since bit.ly/sn-569 was already taken by someone and links to a Google search for "steven gibson smelly old man")

Automated system certificate store checking for Windows workstation and server. Alerts users to the addition and removal of system certificates.

This new, free utility will monitor any changes made to the Windows Certificate Stores on your system. Certificates can be added or removed to your system for a variety of reasons - Windows Updates, new software packages, etc. can make alterations to the certificate store. Unfortunately, some malicious software could also add an "all purpose" certificate and essentially create an attack vector for SSL/TLS man-in-the-middle attacks or provide a foothold for a bad agent to usurp and exfiltrate information from your system without your knowledge. CertWatch performs hourly scans of all system certificate stores and will report any additions or deletions from those stores when changes are made.

20-year-old, designed-in, Windows behavior lets printers install malware.
http://blog.vectranetworks.com/blog/microsoft-windows-printer-wateringhole-attack
http://info.vectranetworks.com/understanding-printer-vulnerabilities
Security researchers with Vectra Threat Labs recently uncovered a critical vulnerability (CVE-2016-3238) which affects ALL VERSIONS OF Microsoft WINDOWS all the way back to Windows 95. The vulnerability is created by the way Windows clients interact with network printers, allowing an attacker to execute code at system level either over a local network or the Internet.

20 years ago Microsoft implemented a VERY DANGEROUS feature known as " Microsoft Web Point-and-Print Protocol" (MS-WPRN) which allows a Windows machine, connecting to a network-hosted printer for the first time, to RECEIVE AND INSTALL A PRINTER DRIVER FROM THE PRINTER.

What could possibly go wrong????

(I am reminded of the very similar Windows Metafile design mistake Microsoft also made during the same time period.)

Vectra Networks: <<quote>>

Most organizations try to apply the principle of least privilege to the devices in their networks. This works pretty well for things like laptops or desktops since the hardware they use doesn't change that often. However printers are a bit different. While they still need drivers, printers need to support virtually any user that wants to connect to them. As end-users move through a build, they naturally want to use the printer closest to them. Mobile users expect to be able to easily connect and use a printer when they come into the office. In addition, most organizations don't standardize on a single printer, and will have multiple models and manufacturers often within a single network.

So instead of having system administrators push all possible printer drivers to all workstations in the network, the solution was to develop a way to deliver the driver to a user device right before the printer is used. And this is where Point-and-Print showed up. This approach stores a shared driver on the printer or print server, and only the users of that printer receive the driver that they need. At first glance, this is a practical and simple solution to driver deployment. The user gets access to the printer driver they need without requiring an administrator – a nice win-win.

The Issue??

The problem is that for this scheme to work nicely from an end-user perspective, an exception was required. Normally, User Account Controls are in place to warn or prevent a user from installing a new driver. To make printing easier, an exception was created to avoid this control. So in the end, we have a mechanism that allows downloading executables from a shared drive, and run them as system on a workstation without generating any warning on the user side. From an attacker perspective, this is almost too good to be true, and of course we had to give it a try.

Researchers at the security firm Vectra Networks discovered that the Windows Print Spooler doesn't authenticate print drivers when installing them from remote locations. That lack of authentication makes it possible for attackers to use several different techniques that deliver maliciously modified drivers instead of the legitimate one provided by the printer maker. The exploit effectively turns printers, printer servers, or potentially any network-connected device masquerading as a printer into an internal drive-by exploit kit that infects machines whenever they connect.

But wait... there's more!! Vectra Networks <<paraphasing a bit for emphasis>>

Infecting Remotely Using Internet Printing Protocol and webPointNPrint
So far we have constrained ourselves to an internal network where a device was either inserted or infected and used to further infect devices connecting to it. Internet printing protocol (IPP) and webpointNprint allow us to extend this issue outside the intranet to the internet. IPP allows for the same mechanism to load drivers from remote, in this case very remote, printers. This can be done with following piece of code from the MS print server....

So… this was "fixed" last Tuesday. What did Microsoft change? They are UNABLE TO CHANGE this behavior without crucially and critically breaking current roaming laptop transparent printer driver installation... So they've added a confirmation dialog. (And we've seen how well those work with, for example, not upgrading to Windows 10.)

It's also possible to disable "Point-and-Print."

Surprising flaw found in Juniper router OS:
https://kb.juniper.net/InfoCenter/index?page=content&id=JSA10755&actp=search
Title: "2016-07 Security Bulletin: Junos: Self-signed certificate with spoofed trusted Issuer CN accepted as valid (CVE-2016-1280)"

Products Affected:
- This issue can affect any product or platform running Junos OS

Problem:
- Junos OS runs PKId for certificate validation. When a peer device presents a self-signed certificate as its end entity certificate with its issuer name matching one of the valid CA certificates enrolled in Junos, the peer certificate validation is skipped and the peer certificate is treated as valid. This may allow an attacker to generate a specially crafted self-signed certificate and bypass certificate validation.

Errata:

- David Redekop (and several others) correct my mistatement
    - SG300's are not iOS (though iOS command-like)

Miscellany

"Three Dumb Routers" ...or... ONE Super-Smart Router!
Ubiquity EdgeRouter X $50, 5-separate interfaces, w/power supply
UPdate the firmware!
Packet switching rate (~500 Mbits)
"Normal Router" is a two-interface NAT connected to a multiport hub or switch. One the LAN side it has a single LOGICAL interface and five to eight physical interfaces.
A "super smart" router has separate network logical interface for every physical interface.
No OpenVPN support… but PPTP and IPsec are both natively supported.
But, everything else is there! (And anyone could run OpenVPN on a Raspberry PI.)

Finished "The Endless North Road"

- I did enjoy it, but it was quite long.
- Now returning to Jack Campbell's Lost Fleet series
- Four books "Beyond the Frontier"

StarTrek premiered in 1966 (when I was 11 years old) and is celebrating it's 50th anniversary.

As Popular Science's coverage put it: "Netflix users worldwide will be able to boldly binge where few have ever binged before. By the end of the year, 727 episodes of all previous Star Trek series will be available worldwide, on Netflix, for streaming.

CBS's much-anticipated new Star Trek TV series, slated for January, will make new episodes available for streaming on Netflix within 24 hours of their debut -- but NOT for viewers in the U.S. and Canada. U.S. viewers will need to go through CBS and Canadian viewers through "Crackle" until the episodes eventually appear on Netflix.

Mr.Robot? Hmmmm....

I hope we soon spend less time inside Elliot's disturbed head and get back to some fun hacking and bringing down Our Evil Corporate Overlords.

Michael Cunningham / @mikecunning

Did you notice the QR code in Elliot's journal actually works and goes to some goofy site?

http://www.confictura industries.com/

http://www.theverge.com/2016/7/14/12187768/qr-code-mr-robot-confictura-industries-usa

This Week in Puzzles / iOS / Android / Steam

"klocki" by Maciej Targoni ($1)

http://appsto.re/us/nOU4bb.i

http://klockigame.com/

Subject: "The future of Hook" / Maciej:

Maciej: I'm thinking constantly about expanding it, but I think I will just go with a new game. The main design idea was to add new game mechanics continually, to maximize "AHA!" moments. At this stage there is not that much interesting things to add. And I don't want to stretch play time just for a sake of it. Game has to be fun from the beginning to the end.

Steve: There is another way to look at it, though... which I will try to explain: Once someone has acquired a new skill or understanding, it can be fun to simply use that new skill, for at least a awhile, even if nothing further is being learned. It can just be a pleasant diversion to solve the puzzle even if nothing more is being learned about the puzzle. Just "working the machine" and accomplishing is enough.

Perhaps the best example is the enduring popularity of crossword puzzles or sodoku. The people who enjoy them are perhaps getting a bit better at them over time, but the pleasure is just in using their brain to work out this particular solution... which always changes.

I wondered whether it would be possible to create a "Hook Level Generator" which is capable of endlessly creating Hook levels... by itself? That could be a 'Hook Pro" or "Hook Ultimate" for $4.99 which people could play and play and play... just for the pure pleasure of solving different Hook puzzles.

Maciej: Its my second "serious" game, so I think I can do better :)
I will stay within puzzle genre and keep exploring it, looking for unique ideas. And my design will keep on going in the direction you are calling a "sweet spot" - slow progress, relaxing, no timer, no rush, no 3-star award system :)

HSF: Common side effects are morning headache and groggy hangover.

- Half dose (750mg) has resolved that for everyone and is still effective.
- I tried half of both and it was still effective.
- Unfortunately I keep adding links to suppliers, and each one sells out:
  Amazon, iHerb, Swanson, Vitamin Shoppe, Drugstore.com & eVitamins -- all sold out.
- Source Naturals caught completely off guard.
- PLEASE do not use another. ONLY that one will work.

SpinRite

Monday, 9:26am / Emett Speer / @BotEmett
I also wanted to add that I am an owner of SpinRite 6.
I was trying to get my last company to purchase a corporate SpinRite 6 license (by purchasing four licenses) for use on the many PC's they own. I demonstrated its abilities by using it to recover a dead RAID 6 array for a customer who was in a panic to get their data back. The company and customer where grateful that I was able to fully recover the full system with no data lost, but they didn't get SpinRite for our PC's.

Messenger -- FaceBook Messenger
Abuse Reporting & Secure Storage Management

"Franking" - Nonrepudiatable message signing.

- Any participant in a secret conversation may voluntarily notify Facebook of abusive content they receive. Facebook uses such reports to identify users who violate Facebook's terms of service. The ability to report abuse does not relax the privacy guarantees inherent in the end-to-end encryption of Secret Conversations. Facebook will never have access to plaintext messages unless one participant in a secret conversation voluntarily reports the conversation.

- The sender of any message must incorporate a "franking tag" which is appended to and encrypted in the message.

- No recipient of a message will display any message without a valid "franking tag" which is validated upon receipt.

- Sender generates a nonce which is (a) appended to the message being sent and (b) keys an HMAC256 of the entire concatenation.

- The 256-bit HMAC output is the "franking tag" for the message.

- Sender destroys the nonce after using it, then the message and the franking tag is forward to Facebook for forwarding to the recipient.

- When Facebook receives the message + franking tag, it uses a secret Facebook key to key another HMAC256 consisting of the concatenation of the franking tag with the "conversation context" (the sender and recipient identitiers and a timestamp). This creates a "Reporting Tag".

- The Reporting Tag, the Franking Tag and the message are then all sent to the recipient.

- Upon receiving report of abuse, the recipient returns the:
    - Full decrypted message plaintext.
    - Reporting Tag, the Nonce appended to the end of the plaintext
    - The conversation context.

- Face recomputes the original franking tag, then using its Facebook key, recomputes the Reporting tag. Only if *everything* is true: the endpoint user identifiers, the "when" timestamp, and the unaltered message content, will everything match... and there is no way the sender can claim that they didn't send what they sent, when they sent it.

Secret Conversation Secure Storage

- Secret Conversations plaintext messages are stored permanently only on the devices that participate in each conversation.

- Plaintext messages are protected using on-device symmetric-key encryption and optional Disappearing Messages functionality.

- On-device encryption ensures that messages stored permanently on a particular device are only accessible while a user is authenticated to Facebook.

- The operational requirements are:
  - Messenger allows users to switch Facebook accounts.
  - While a second user is logged in to a particular device, messages of the first user are not accessible.
  - However, when the first user returns to the same device they will find their messages intact.

- To achieve these requirements, clients employ two encryption keys: Local Key and Remote Key.

- Both these keys are used for AES-GCM encryption.

- Local Key is generated on-device and never leaves the device it was generated on. It is used to encrypt plaintext messages before these are stored permanently on a device.

- Remote Key is a long-term, user-specific key held on Facebook and delivered to the device when a user authenticates. The Remote Key is used to encrypt the Local Key in local storage. When Messenger switches accounts, the device persists an encrypted version of the Local Key and erases Remote Key.

- Upon successful reauthentication, the device obtains Remote Key from Facebook, uses it to decrypt the Local Key and regain access to messages.

Disappearing Messages Feature:

- Disappearing Messages ensures that messages are no longer visible within a selected time after they are sent or received.

- Disappearing Messages incorporate a timeout which is added to the message data structure before serialization and encryption.

- Both devices honor and automatically hide messages that specify a timeout once the timeout has elapsed.

- However, the actual physical deletion of message plaintext from local storage does not occur immediately, but does occurs shortly after each message has expired, in order to enable abuse reporting in the interim.

CryptoDrop

CryptoDrop: Researchers develop a way to stop ransomware
http://www.cise.ufl.edu/~traynor/papers/scaife-icdcs16.pdf

Goal: Develop an early warning system for ransomware

- Identify three primary indicators suited to detect malicious file changes.
- Perform the most extensive analysis of encrypting ransomware to date.

Indicators:

File Type Changes:

- Files of known types have well-known structures, especially in their headers at the top of the file.

- The "file" utility is a popular program for determining file type. The default “magic” database library contains hundreds of file type signatures, ranging from specific programs (“Microsoft Word 2007+”) to general content (“Unicode text, UTF-7”). They use this tool to track the file type both before and after a file is written. If this type changes, they can infer that some transformation has occurred.

Similarity Measurement:

- Strong encryption should produce output that provides no information about the plaintext content. They assume that the output of ransomware-encrypted user data is completely dissimilar to its original content.

- The use of similarity-preserving hash functions (sdhash). These hashes differ from traditional hash functions because they contain some information about the source file in their output. Through measuring the similarity of two versions of the same file, they gain information about dissimilarity.

- http://roussev.net/sdhash/tutorial/03-quick.html

- sdhash outputs a "comparison metric" 0-100 for the similarity of two files.

- Given the similarity hash of the previous version of a file, a comparison with the hash of the encrypted version of that file should yield no match, since the ciphertext should be indistinguishable from random data.

Shannon Entropy

- This one is quite straightforward (though compression would fool it).
- Either encrypted *or* compressed files would be expected to have very HIGH entropy.

Testing in the lab:

- We obtained 2,663 malware samples from VirusTotal using ransomware-related search terms and known variant names.

- We confirm in this data that there is a relative lack of diversity within each family's behavior. In both the TeslaCrypt and CTB-Locker families, which consist of over half of the working samples we obtained, two or fewer samples showed behaviors beyond their families' primary behavior class. The greatest diversity was found in the Filecoder and CryptoLocker families, though these two family names are often used as generic ransomware detection names.

- Due to the homogeneous nature of the behavior in each sample, our data shows that CryptoDrop remains robust against many forms of encrypting ransomware despite low counts of usable samples in some families. Because our study covered nearly four times the number of families of the previous study and there was little diversity within families, there was little need to collect additional samples.

- CryptoDrop detected the behavior of all 492 widely differing samples, quickly protecting the majority of victim's data with as few as zero files encrypted before detection. This result highlights the required actions of ransomware, and the effectiveness of our indicators at detecting this type of malware. Below, we discuss the ability of our system to protect user data and the effectiveness of union indication.

- File Format attack frequency:
    - PDF, ODT, DOCX, PPTX, TXT, MOV, ZIP, MD...

False Positives:

- While CryptoDrop is effective at quickly detecting ransomware, we note that any evaluation of its real-world utility must also include a discussion on incorrect detection of benign activity. False positive analysis for a system such as CryptoDrop is challenging since its analysis requires changes to be made to a user's protected documents. Techniques used in static malware analysis works (e.g., providing a set of known-good binaries alongside bad ones) will not work since CryptoDrop does not analyze binaries for malicious traits. Likewise, techniques used in dynamic malware analysis (e.g., passively observing benign activity on a system and running the detector on it later) will not work since CryptoDrop needs to measure the user's documents before and after each change. Below, we discuss the results of our experiments with benign programs and show that our system remains robust.
- We evaluated thirty common Windows applications on the same virtual machine configuration used to test malware samples and found only ONE false positive. That false positive, 7-zip, was expected, as it reads a large number of disparate files and generates high entropy output (similar to ransomware).

Riffle

MIT anonymity network promises to be more secure than Tor
https://www.engadget.com/2016/07/11/mit-anonymity-network-more-secure-than-tor/
https://news.bitcoin.com/mits-riffle-claims-anonymous-tor/
http://dspace.mit.edu/bitstream/handle/1721.1/99859/927718269-MIT.pdf?sequence=1

Improves the efficiency of a fundamentally inefficient solution to solve a practically impossible problem.