Security Now! #586 - 11-15-16

The BlackNurse Attack

This week on Security Now!

- Results from our listener's informal CAIDA spoofing testing,
- LessPass turned out to be even less than it appeared,
- My day at Yubico,
- A bunch more IoT news
- News from PwnFest & Mobile Pwn2Own
- A bit of miscellany, including the probable elimination of Dark Matter,
- A new WiFi field disturbance attack,
- A whacky Kickstarter "fingerprint" glove,
- And the "BlackNurse" reduced-bandwidth DoS attack.

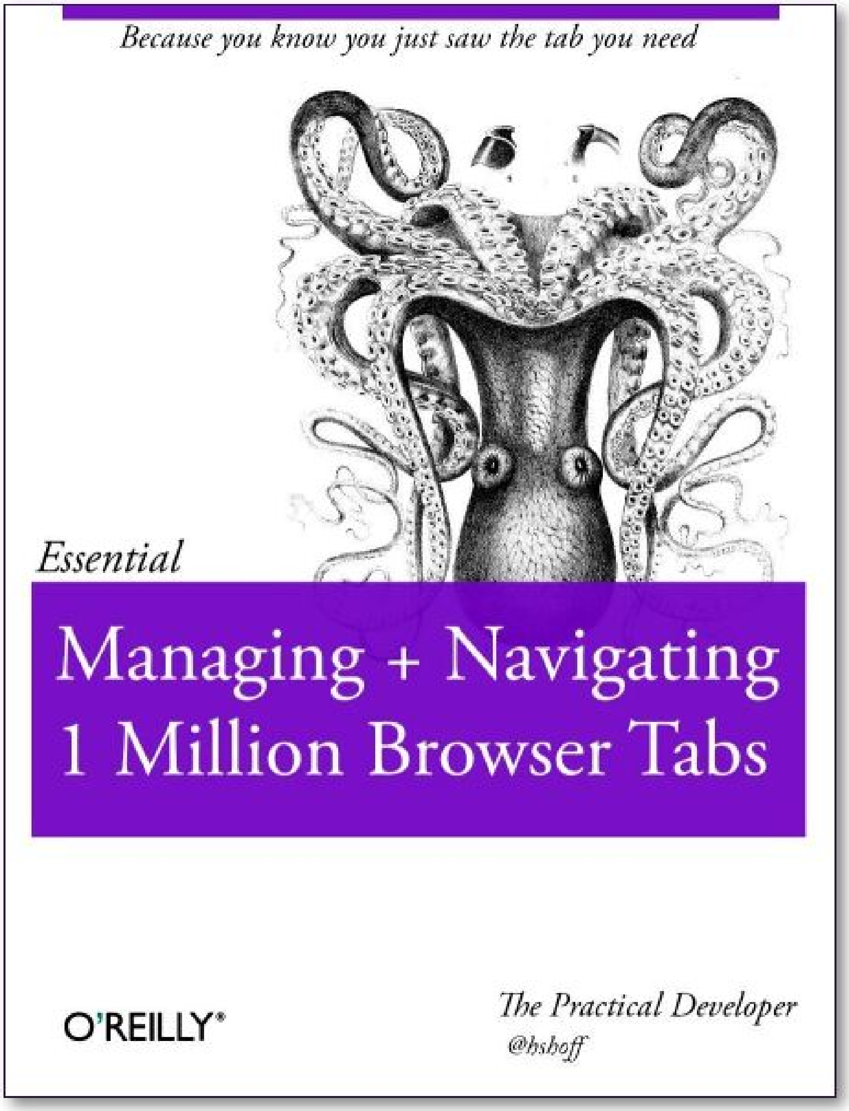
Because you know you just saw the tab you need
Essential
Managing + Navigating 1 Million Browser Tabs
O'REILLY®
The Practical Developer
@hshoff

Security News

CAIDA IP Spoofing Test Results

- Thanks to the hundred of results shared.
- Only ONE router didn't check the low-order IP address byte:
    - Netgear N600 WNDR3700v3 w/current firmware.

Revisiting LessPass

- The guy who did "LessPass" shot me a tweet:
  @SGgrc "Hope he doesn't attempt to obtain a trademark on that" no we are not trying to obtain a trademark on LessPass.

And... one of listeners named Adam took a look at the code... and was horrified by what he found.

- Adam reminded me that I had said that I hadn't bothered to look at the code since it was now trivial to do this correctly, and I assumed that this guy had.

- But based upon Adam's analysis, nothing could be further from the truth.

- So if my discussion last week of the many obvious UI reasons why LessPass is not practically useful didn't dissuade you, the following should:

- Here's how to solve this problem:
    - Hash followed by successive long division by the modulus of the chosen alphabet size.
    - The remainder of each division will be 0 to one less than the alphabet size.
    - Map that one-to-one into the chosen character.

- How to solve the problem horribly:
    - Step through the selected check boxes sequentially, using some bits from the source hash output to choose from among the characters in that set.
    - This will, of course, result in absolutely predictable sequential ordering of characters -- lowercase alpha, uppercase alpha, number, special character... repeat.
        - (Adam notes that the lowercase and uppercase classes are, for some reason, each divided into vowels and consonants, making this even worse... since the range of possibilities for those character positions becomes much smaller and easier to predict.)
    - And depending upon how he mapped the bits of the source binary to the characters, which I no longer have ANY confidence in, it might well result in a non-uniform distribution of characters within each sub-set.

- Anyone simply observing several outputs of this thing would have quickly noticed the troubling pattern... but then so would any attacker.

- So we wouldn't have even needed open source to know that something was VERY wrong with it.

- Thank you, Adam, for taking time to look under the covers, and to let us know what you found.

- This should be a cautionary warning to us: No matter how well intentioned, not everyone is equally capable of solving important problems correctly. (And even those who are can still make mistakes.) But the last thing anyone should trust is ad-hoc constructions such as this, or as another example, the bizarre "encryption" used by Telegram.

- There are well known and increasingly time proven the tested ways to solve all of these problems... and there's no good reason to make up new stuff "just because".

- So... on balance, I'd have to say that this thing called "LessPass" was more much appropriately named than I thought.

My day at Yubico with Stina and Jakob

- Last Thursday
- 38-page "Full Feature Walkthrough" presentation.
    - Will be finalized and posted.

Joe Rodricks (@Joerodricks)

- @SGgrc @WIRED why is there a silent audio track playing when I visit wired.com on my iPhone? Not cool.
- I resume this is a technique for monitoring how long the user remains on the page.
- The trouble is... Wired may have ample bandwidth, purchasing it in bulk at massive discount... but their visitors may have fixed data rate plans which are decidedly less than unlimited.
- There seems to be a general diminishing concern for website visitors...

Spotify is writing massive amounts of junk data to storage drives

- http://arstechnica.com/information-technology/2016/11/for-five-months-spotify-has-badly-abused-users-storage-drives/
- Dan Goodin: "Streaming app used by 40 million writes hundreds of gigabytes per day."
- QUOTE: For almost five months—possibly longer—the Spotify music streaming app has been assaulting users' storage devices with enough data to potentially take years off their expected lifespans. Reports of tens or in some cases hundreds of gigabytes being written in an hour are not uncommon, and occasionally the recorded amounts are measured in terabytes. The overload happens even when Spotify is idle and isn't storing any songs locally.
- Dan writes: The behavior poses an unnecessary burden on users' storage devices, particularly solid state drives, which come with a finite amount of write capacity. Continuously writing hundreds of gigabytes of needless data to a drive every day for months or years on end has the potential to cause an SSD to die years earlier than it otherwise would. And yet, Spotify apps for Windows, Mac, and Linux have engaged in this data assault since at least the middle of June, when multiple users reported the problem in the company's official support forum.

- Three Ars reporters who ran Spotify on Macs and PCs had no trouble reproducing this effect which had been reported not only on the Spotify forum but also on Reddit, Hacker News, and elsewhere. The Spotify app wrote from 5 to 10 GB of data in less than an hour on Ars reporters' machines, even when the app was idle. Leaving Spotify running for periods longer than a day resulted in amounts as high as 700 GB.

- The problem is reported to have been fixed in v1.0.42 which is in the process of being rolled out.

- The bug appears to be that the SQLite database is being repeatedly and redundantly compacted.
    - VACUUM: "The VACUUM command rebuilds the database file, repacking it into a minimal amount of disk space."

All your lightbulbs are belong to us!

- BoingBoing: A lightbulb worm could take over every smart light in a city in minutes
- https://boingboing.net/2016/11/09/a-lightbulb-worm-could-take-ov.html

- *Cory Doctorow:*
  Researchers from Dalhousie University (in Canada) and the Weizmann Institute of Science (in Israel) have published a working paper detailing a proof-of-concept attack on smart lightbulbs that allows them to wirelessly take over the bulbs from up to 400m, write a new operating system to them, and then cause the infected bulbs to spread the attack to all the vulnerable bulbs in reach, until an entire city is infected.

  The researchers demonstrate attacking bulbs by drone or ground station. The demo attacks Philips Hue lightbulbs, the most popular smart lighting system in the market today.

  Philips Hue use Zigbee for networking. Zigbee is a wireless protocol designed for low-powered Internet of Things devices, and it has many built-in security features. The most important of these is that once a device is initialized as part of a Zigbee network, it can't be hijacked onto a rival network unless you can bring a controller into close proximity to it (a couple centimeters away). However, there is a fatal flaw in the Zigbee implementation in the Hue system, and the researchers showed that they could hijack the bulbs from nearly half a kilometer away (this attack is only possible because Zigbee doesn't encrypt all traffic between devices).

  The Hue system also has safeguards to prevent malicious tampering: updates have to be cryptographically signed using a very strong algorithm or they will be rejected by Hue systems. The researchers were easily able to extract the signing keys -- which are the same for all Philips Zigbee products -- and use them to sign their own malicious updates.

  Thus armed, the researchers were able to take over any Philips Hue system.

  There are many ways that a hijacked Hue system can be used to cause mischief. Zigbee uses the same radio spectrum as wifi, so a large mesh of compromised Zigbees could simply generate enough radio noise to jam all the wifi in a city. Attackers could also brick all the Hue devices citywide. They could use a kind of blinking morse code to transmit data stolen from users' networks. They could even induce seizures in people with photosensitive epilepsy.

The fact that the attack targets devices by Zigbee signals -- rather than over the internet -- means that it is virtually impossible to defend against through traditional methods like firewalls.

- Not so long ago this scenario would have been seen as far fetched at best. Now?? Its exploitation is a virtual certainty.

The "Web Of Trust" browser extension seems less than trustworthy, itself.

- http://www.pcmag.com/news/349328/web-of-trust-browser-extension-cannot-be-trusted

- I'm a huge fan of browser extensions. They have become an essential part of my daily management of my portal to the Internet. They add security, block annoyances, add extra features, and act as a testing ground for functionality that may eventually become a standard part of the most popular browsers.

- Against the background of this nice ecosystem, the German TV channel NDR has uncovered a serious breach of privacy by the Web Of Trust (WOT) service, which over 140 million Web surfers trust to help keep them safe online. Web-Of-Trust has been around since 2007 and described itself as a "Safe Web Search & Browsing" service. This boils down to a crowdsource-driven website reputation and review system. Users may view ratings on a per-site basis for trustworthiness or child safety and add their own ratings.

- The NDR investigation discovered that while you have the WOT extension installed, extensive data collection is going on in the background. But WOT not only collects and records data on a per-user basis, it then analyzes and sells it to third parties.

- The WOT Privacy Policy states that your IP, geographic location, device type, operating system, browser, the date and time, Web addresses, and overall browser usage are all collected, but that it is "non-identifiable." But NDR found that it was a simple task to link the anonymized data to individual users of the service. The data retrieved included:

    - Account name
    - Mail address
    - Travel plans
    - Illnesses
    - Sexual preference
    - Drug consumption
    - Confidential company information
    - Ongoing police investigations
    - Browser surfing activity including all sites visited

- NDR pulled all of this information from a small data sample of around 50 users. But WOT has access to data for all 140 million users.

- Mozilla immediately removed the WOT extension from its Firefox Add-ons page due to guideline violations. It seems likely other software that supports WOT will follow.

- Anyone currently using the WOT extension should seriously consider whether they wish to continue doing so. WOT also has a mobile app, which won't be immune to this data collection.

Adult Friend Finder -- Hacked. (Reported Sunday)

- https://www.leakedsource.com/blog/friendfinder
- http://www.csoonline.com/article/3132533/security/researcher-says-adult-friend-finder-vulnerable-to-file-inclusion-vulnerabilities.html
- "LeakedSource", which follows and aggregates site breaches sensationalized the breach on Adult Friend Finder with the headline: "Sexual secrets for hundreds of millions exposed in largest hack of 2016"

- Friend Finder Network Inc is a company that operates a wide range of 18+ services and was hacked in October of 2016 for over 400 million accounts representing 20 years of customer data which makes it by far the largest breach we have ever seen -- MySpace gets 2nd place at 360 million. This event also marks the second time Friend Finder has been breached in two years, the first being around May of 2015.

- AdultFriendFinder.com
  103,070,536 passwords already plainly visible
  232,137,460 passwords hashed with SHA1
  99.3% of all passwords from this website are now plaintext (cracked).
- Cams.com
  21,422,277 passwords already plainly visible
  41,209,412 passwords hashed with SHA1
  96.8% of all passwords from this website are now plaintext (cracked).
- Penthouse.com
  495,720 passwords already plainly visible
  6,678,239 passwords hashed with SHA1
  99.9% of all passwords from this website are now plaintext (cracked).
- Stripshow.com
  342,889 passwords already plainly visible
  1,080,303 passwords hashed with SHA1
  99.95% of all passwords from this website are now plaintext (cracked).
- iCams.com
  272,409 passwords already plainly visible
  863,317 passwords hashed with SHA1
  99.96% of all passwords from this website are now plaintext (cracked).

- Total: 99.0% of all available passwords are now visible in plaintext

- Many of the site's user passwords are posted... and most are quite NSFW.

A turf war over IoT devices

- Competing hackers dampen the power of Mirai botnets
- http://www.pcworld.com/article/3138040/security/competing-hackers-dampen-the-power-of-mirai-botnets.html

- DYN was hit with a massive attack by the IoT-based Mirai botnet.

- But the release of Mirai's source code two months ago, in September, quickly resulted in huge competition for the limited pool of infectable devices, and thus a massive fracturing into as many as 52 significantly smaller botnets. While this permits many more simultaneous attacks, unless they were to team up, their individual attack scales would be reduced in size due to a lack of number.

- Remember that Mirai only lives in RAM... so the reset/reboot of any Mirai-hosting device flushes the bot that had been living there and reopens that device to reinfection by the first bot that comes along.

The "S" in IoT is for Security!

PwnFest 2016:

- Google Pixel Phone and Microsoft Edge Hacked at PwnFest 2016
- "Google Pixel Gets Hacked In Under A Minute"
  http://www.ubergizmo.com/2016/11/google-pixel-gets-hacked-in-under-a-minute/
- "Google's new Pixel phone hacked in 60 seconds"
  http://mashable.com/2016/11/12/google-pixel-hacked-60-seconds/#CY2Q6Y92ASqA

- In my opinion, these attention-grabbing headlines missed the point... but more on that in a second.

- Earlier this month, Adrian Ludwig, the director of Android security at Google, told Motherboard at a security conference that the Pixels are as secure as iPhones. He said: "For almost all threat models [the Pixels and iPhone] are nearly identical in terms of their platform-level capabilities."

- He may well wish to believe that, but we all know that for security the proof is both in the design and in the implementation... which only history can judge. And so far, history has not been kind to Android. As we know, Android has traded openness for security... and it struggling to offer both. But today it is nowhere nearly as secure -- in the field -- as the iPhone.

- And a group of Chinese white hat hackers hacked a brand new Google Android Pixel in 60 seconds last Friday at the PwnFest hacking competition that took place in Seoul on Friday, according to The Register.
- The hackers, who work at Qihoo 360, a security solutions company, won $120,000 in cash after demonstrating an exploit that cracked open the Android and gave them full remote access as well as access to personal information such as messages, phone calls, contacts

and photos.

- Anyone can make a mistake, and Google is playing security catch up. But what they CAN and SHOULD be proud of is that they had the newly discovered problem patched within 24 hours!

- Oh... and they also closed a previous hole that enabled rival white hat hackers at Tencent's Keen Labs to also <<ahem>> breach the Pixel's security at the Mobile Pwn2Own event in Japan last month.

Shove this in your pipe and smoke it!

- Adding a bit more security to the CURL | SH hack
- https://github.com/awalGarg/curl-tap-sh
- The concerns over "curl foo/bar | sh"
- TAP first collects all of the data from curl, saves it to a temp file, opens that file in your $EDITOR (or vim if not specified), for review it.
- You can make changes to it if you want. If you write the file and close the editor successfully (i.e., the editor returns exit code 0), then TAP sends the saved output (including your edits, if any) along the pipe.
- This also shields against a timing attack which detects curl | sh server-side.

OpenSSL has been patched again (for a high severity vulnerability).

- http://thehackernews.com/2016/11/openssl-patch-update.html
- https://www.openssl.org/news/secadv/20161110.txt
- Three vulnerabilities fixed, High, Medium, and Low.
- Only affecting users of v1.1.0 (not prior to v1.1.0
- Severity: High
  "TLS connections using *-CHACHA20-POLY1305 ciphersuites are susceptible to a DoS attack by corrupting larger payloads. This can result in an OpenSSL crash. This issue is not considered to be exploitable beyond a DoS.

- OpenSSL 1.1.0 users should upgrade to 1.1.0c

- This issue was reported to OpenSSL on 25th September 2016 by Robert Swiecki (Google Security Team), and was found using honggfuzz. The fix was developed by Richard Levitte of the OpenSSL development team.

Eugene Kaspersky is unhappy with Microsoft

Kaspersky accuses Microsoft of anticompetitive bundling of antivirus software

- http://arstechnica.com/information-technology/2016/11/kaspersky-accuses-microsoft-of-anticompetitive-bundling-of-antivirus-software/
- https://eugene.kaspersky.com/2016/11/10/thats-it-ive-had-enough/

Casting themselves in the role of David to Microsoft's Goliath.

Reading through Eugene's long list of complaints I was immediately put in mind of the Get Windows 10 debacle.

Eugene writes:

- Several years ago Microsoft decided to overhaul the Windows platform. Ostensibly this was in the name of better ease of usage, security, performance and so on. Behind the scenes what Microsoft was up to was elegantly seizing niche markets: squeezing independent developers out of them, taking their place, and offering users their own products, which in many cases were in no way better.

- Users of Windows 10 have been complaining that the system is changing settings, uninstalling user-installed apps, and replacing them with standard Microsoft ones. A similar thing's been happening with security products.

- When you upgrade to Windows 10, Microsoft automatically and without any warning deactivates all 'incompatible' security software and in its place installs… you guessed it – its own Defender antivirus. But what did it expect when independent developers were given all of one week before the release of the new version of the OS to make their software compatible?

- Even if users have compatible protection from an independent developer already installed, Defender appears with an alarming window. It fairly shouts that Defender is switched off, because you've some other AV installed. There's a big juicy Defender 'Turn on' button too. Of course, many users will be inclined to press this button: 'well, it's from Microsoft – the people who make the OS; must be good; no harm in turning it on for sure'. But pressing the big juicy button also deactivates your existing AV. (A user only finds this out from a tiny text in a pop-up window (and they need to know how to get that window to pop-up):

- Microsoft has even limited the possibility of independent developers to warn users about their licenses expiring in the first three days after expiration. Actually, a warning is there, but it's buried in a Windows Security Center notification, which hardly ever gets read. What's the big deal about three days? It's a big deal because this is the crucial period during which a significant number of users seek extensions of their security software licenses. And if a user forgets to renew a license, then Microsoft deactivates the existing AV, and turns on Defender.

- Microsoft has introduced a limit on the number of antiviruses you can have on a PC: one (or two – if one of them is Defender; see below). At first glance this looks like sense: all for a more comfortable user experience. But the devil’s in the details: Let’s say you’ve an independent AV. You intentionally – or not (e.g., with bundled software) – install a trial version of a different AV, but forget to delete it or don't purchase a license for it. When the trial period is up, Windows quietly turns off both AVs, and turns on Defender!

In a presentation titled: "Windows 10 - Protecting Device Integrity" the presenter says: “I want you to think about kicking out the third party antivirus because we’ve got a great solution right now and it’s going to be even better in the months to come.”

But Eugene misses the significant benefits created for them:

- Being a "monopoly OS" meant that they needed to support many fewer platforms to get good market penetrating coverage.

And Eugene has apparently missed seeing the writing on the wall.

- This is what Microsoft does... and, to their credit, they generally do it gently.
- Once upon a time there was a huge industry for Windows add-on firewalls.
- Then XP added a built-in firewall. But it didn't turn it on by default for several years.

And I've been in that position myself:

- I'll never forget the dinner I shared with "The brads" (Silverberg and Chase)
- ScanDisk in v6.2... to replace CHKDSK.
- For years we were having to answer the question: “How is SpinRite better than ScanDisk.”
- But many people never asked… they just used ScanDisk and then gave up when it didn’t restore the health of the hard drive and its data.

Miscellany

WestWorld : Renewed for a second season

- A delightful 60 minutes every Sunday.

Arrival : A very nicely done First Contact movie

- It's "Close Encounters of the 3rd Kind" meets "Slaughterhouse Five".

A new theory of gravity might solve the dark matter conundrum.

- http://m.phys.org/news/2016-11-theory-gravity-dark.html
- So what's the problem?
  - We think we know how gravity works, even though its theory has always been at odds with quantum mechanics.
  - We've always known that Newton's classic theory of gravity and quantum mechanics cannot both be correct.
  - The trouble is... observed galaxies are spinning faster than they should be able to, without flying apart.
  - The only way to explain this is for them to contain more gravity than we would observationally predict. So this has forced the assumption that there must be some

so-called "Dark Matter", so named because it must exist for its gravitational sake, while remaining unseen.
    - But there have also been two really annoying problem with this:
        - We need a LOT of it -- such that about 80% of the universe's matter would be dark.
        - And... we've been looking for it, searching for dark particles, anything... and found nothing.

- So, nearly seven year ago, back at the end of 2009, a Dutch theoretical physicist, Erik Peter Verlinde, at the University of Amsterdam's Institute for Theoretical Physics, introduced a theory known as "entropic gravity." According to this theory, "gravity exists because of a difference in concentration of information in the empty space between two masses and its surroundings; he also extrapolates this to general relativity and quantum mechanics. He said in an interview at the time: "On the smallest level, Newton's laws don't apply, but they do for apples and planets. You can compare this to pressure of gas. Molecules themselves don't have any pressure, but a container of gas does."

  In his 2010 article (On the origin of gravity and the laws of Newton), Verlinde showed how Newton's famous second law, which describes how apples fall from trees and satellites stay in orbit, can be derived from these underlying microscopic building blocks. Extending his previous work and work done by others, Verlinde now shows how to understand the curious behavior of stars in galaxies without adding ANY of the presumed (but still missing) dark matter.

- Last Tuesday, on November 8th, Verlinde published a new paper showing how his theory of gravity accurately predicts the velocities by which the stars rotate around the center of the Milky Way, as well as the motion of stars inside other galaxies.

- Verlinde writes: "We have evidence that this new view of gravity actually agrees with the observations. At large scales, it seems, gravity just doesn't behave the way Einstein's theory predicts."

- At first glance, Verlinde's theory presents features similar to modified theories of gravity like MOND (MOdified Newtonian Dynamics (1983)). However, where MOND tunes the theory to match the observations, Verlinde's theory starts from first principles and apparently arrives at correct results without any tweaking.

Robot 'sets new Rubik's Cube record'

- http://www.bbc.com/news/technology-37925028

- At the Electronica tradeshow in Munich.

- German developer/manufacturer "Infineon" promoting automated driving subsystems which, they say, need to offer "very low latencies and absolutely reliable and quick technologies."
- The official Rubik's Cube record for a human is 4.904 seconds, which was set by a 14-year-old boy in 2015. And in recent days another youth completed the task in 4.74

seconds.

- A vision-driven mechanical Rubik's cube solver, using a special friction-reduced cube for speed.
- This was approved by the "World Cube Association"
- The machine's camera were shuttered, and the cube scrambled without the knowledge of the machine.
- 0.637 seconds to solve a cube in 21 rotations.

SpinRite

Paul O Kirwan in the UAE
Subject: SpinRite in Riyadh, the Worlds Lagrest Airport Project

Hi Steve, Just a testimonial and a thank you for your product.

I first came across it back in the 1980's when we were opening Riyadh Airport. At that time we had a small selection of PC's for VIP (Ruling Family and Air force Generals) staff, and back then the drives were very unreliable. A combination of heat , dust, cigarette smoke (they all smoked heavily) meant it was a full time job to keep these things working. And nothing was able to do that except SpinRite.

I now own my own copy and, believe me, it has saved my bacon more than a few times. I have used it many times since, and now for preventative maintenance. I talk about it because it's often amazing, and people are skeptical when I tell them that a 30+ year old piece of software can recover their drive... but they are always impressed when it does.

Thanks again. The interface is so familiar I can almost run it blindfolded!

Thanks to you and Leo for a great podcast and thank you for such a great and enduring product.

Coming up...

- A new, local WiFi field disturbance side-channel attack,
- Synthetic Fingerprint Gloves on kickstarter,
- And "The BlackNurse Attack" which allows a single attacking host with just a 15 mbps connection to hold a website offline.

SPONSOR BREAK

Wi-Fi Signal Interference Can Leak Your Passwords and Keystrokes
http://thehackernews.com/2016/11/hack-wifi-password.html
http://dl.acm.org/citation.cfm?id=2978397
http://dl.acm.org/ft_gateway.cfm?id=2978397&ftid=1805782&dwn=1&CFID=864955056&CFTOKEN=28277939

<ABSTRACT> In this study, we present WindTalker, a novel and practical keystroke inference framework that allows an attacker to infer the sensitive keystrokes on a mobile device through WiFi-based side-channel information. WindTalker is motivated from the observation that keystrokes on mobile devices will lead to different hand coverage and the finger motions, which will introduce a unique interference to the multi-path signals and can be reflected by the channel state information (CSI).

The adversary can exploit the strong correlation between the CSI fluctuation and the keystrokes, to infer the user's numeric input. WindTalker presents a novel approach to collect the target's CSI data by deploying a public WiFi hotspot. Compared with the previous keystroke inference approach, WindTalker neither deploys external devices close to the target device nor compromises the target device. Instead, it utilizes the public WiFi to collect user's CSI data, which is easy-to-deploy and difficult-to-detect.

(Note, side channel attacks, being typically passive, are almost always impossible to attack.)

In addition, it jointly analyzes the traffic and the CSI to launch the keystroke inference only for the sensitive period where password entering occurs. WindTalker can be launched without the requirement of visually seeing the smart phone user's input process, backside motion, or installing any malware on the tablet. We implemented Windtalker on several mobile phones and performed a detailed case study to evaluate the practicality of the password inference towards Alipay, the largest mobile payment platform in the world. The evaluation results show that the attacker can recover the key with a high successful rate.

CSI:

- Multiple antenna MIMO systems use Channel State Information, which is essentially signal metadata, to dynamically improve the performance of the system and lower its bit error rate.
- Wikipedia defines CSI this way: "In wireless communications, channel state information (CSI) refers to known channel properties of a communication link. This information describes how a signal propagates from the transmitter to the receiver and represents the combined effect of, for example, scattering, fading, and power decay with distance. The method is called Channel estimation. The CSI makes it possible to adapt transmissions to current channel conditions, which is crucial for achieving reliable communication with high data rates in multiantenna systems."
- So, in other words, in order to pull off the truly remarkable performance of today's WiFi systems (which is so easy for us to take for granted because we see them as black boxes) it was necessary to add a sophisticated supervisory layer which rides above and continuously monitors the radio link. The instantaneous, digitized, CSI information is

available through the radio link API... and, as this paper shows, can be used to infer aspects of the local transmitting environment... to create a side-channel attack on WiFi connected smartphones.

The headline read: “**These Synthetic Fingerprint Gloves Can Unlock Your Phone”**

- http://www.popularmechanics.com/technology/design/a23860/synthetic-fingerprint-gloves/
- A company called Nanotips thinks it can solve the annoying problem of removing gloves to access your fingerprint-sensor locked smartphone in the winter. The product, Taps, is surprisingly lo-fi. It's essentially a fingerprint-shaped sticker made of military-grade polyurethane that you can stick onto the end of any glove. This fingerprint isn't yours--it's a synthetic, individual fingerprint that you can train your phone to recognize the same way it would your own. These synthetic fingerprints don't actually look like fingerprints, but they function the same way as any real finger, creating a recognizable pattern that your phone can remember and use in the future. Most apps allow at least one or two fingerprints to serve as keys, so Taps merely replaces one of your own fingerprints as a biometric password.

  Of course, this technology presents a massive security problem--suddenly, anyone with your glove can access anything you have locked this way on your phone. This design flaw eliminates most of the usefulness of Taps, as the reason to use a fingerprint key in the first place is generally that it's more secure than a password as its unique to you. No one can steal your fingerprints, but they could definitely steal a glove with a synthetic fingerprint stuck to it.

  However, convenience is a great motivator. Taps is currently on Kickstarter where it has raised $2,000 over its $5,000 goal. A synthetic fingerprint of your own can be ordered for just $8, and ships next month. Keep an eye out for an increase in glove theft.

- Now at $11,027 pledged of $5,907 goal.

- TAPS - Touchscreen Sticker w/ Touch ID. Ships Before X-mas https://www.kickstarter.com/projects/nanotips/taps-touchscreen-sticker-w-touch-id-ships-before-x

The "BlackNurse" low bandwidth "packet rate" attack.

- http://soc.tdc.dk/blacknurse/blacknurse.pdf

- Even A Single Computer Can Take Down Big Servers Using BlackNurse Attack
  - http://thehackernews.com/2016/11/dos-attack-server-firewall.html

- Bandwidth flood versus Packet-Rate Attack.
- Switching Rate Optimization:
  - Fast path vs Slow path.
  - "Ping" ICMP Echo request and reply can be easily handled by automation.
  - But some notifications need to be passed up for higher level management.

- ICMP: Destination Unreachable / Port Unreachable

- RFC 1122: A host SHOULD generate Destination Unreachable messages with code: 2 (Protocol Unreachable), when the designated transport protocol is not supported; or 3 (Port Unreachable), when the designated transport protocol (e.g., UDP) is unable to demultiplex the datagram but has no protocol mechanism to inform the sender.

- A Destination Unreachable message that is received MUST be reported to the transport layer. A transport protocol that has its own mechanism for notifying the sender that a port is unreachable (e.g., TCP, which sends RST segments) MUST nevertheless accept an ICMP Port Unreachable for the same purpose.