# Security Now! #692 - 12-04-18

## “GPU RAM Image Leakage”

### This week on Security Now!

This week we discuss another Lenovo SuperFish-style local security certificate screw up, several new, large and high-profile secure breach incidents... and what they mean for us, the inevitable evolution of exploitation of publicly exposed UPnP router services, the emergence of "Printer Spam", how well does ransomware pay? We have an idea now. The story of two iOS scam apps and a false positive Bing warning, progress on the DNS over HTTPS front, and rumors that Microsoft is abandoning their EdgeHTML engine in favor of Chromium. We also have a bit of miscellany, new of a Cyber Security related Humble Book Bundle just in time for Christmas, a bit of closing the loop feedback, and then we'll discuss some new research that reveals that it's possible to recover pieces of web browser page images that have been previously viewed.

This week’s Security Now! Picture of the Week
Needed a page all to itself, so…

“Shelves?  Who needs Shelves?”...

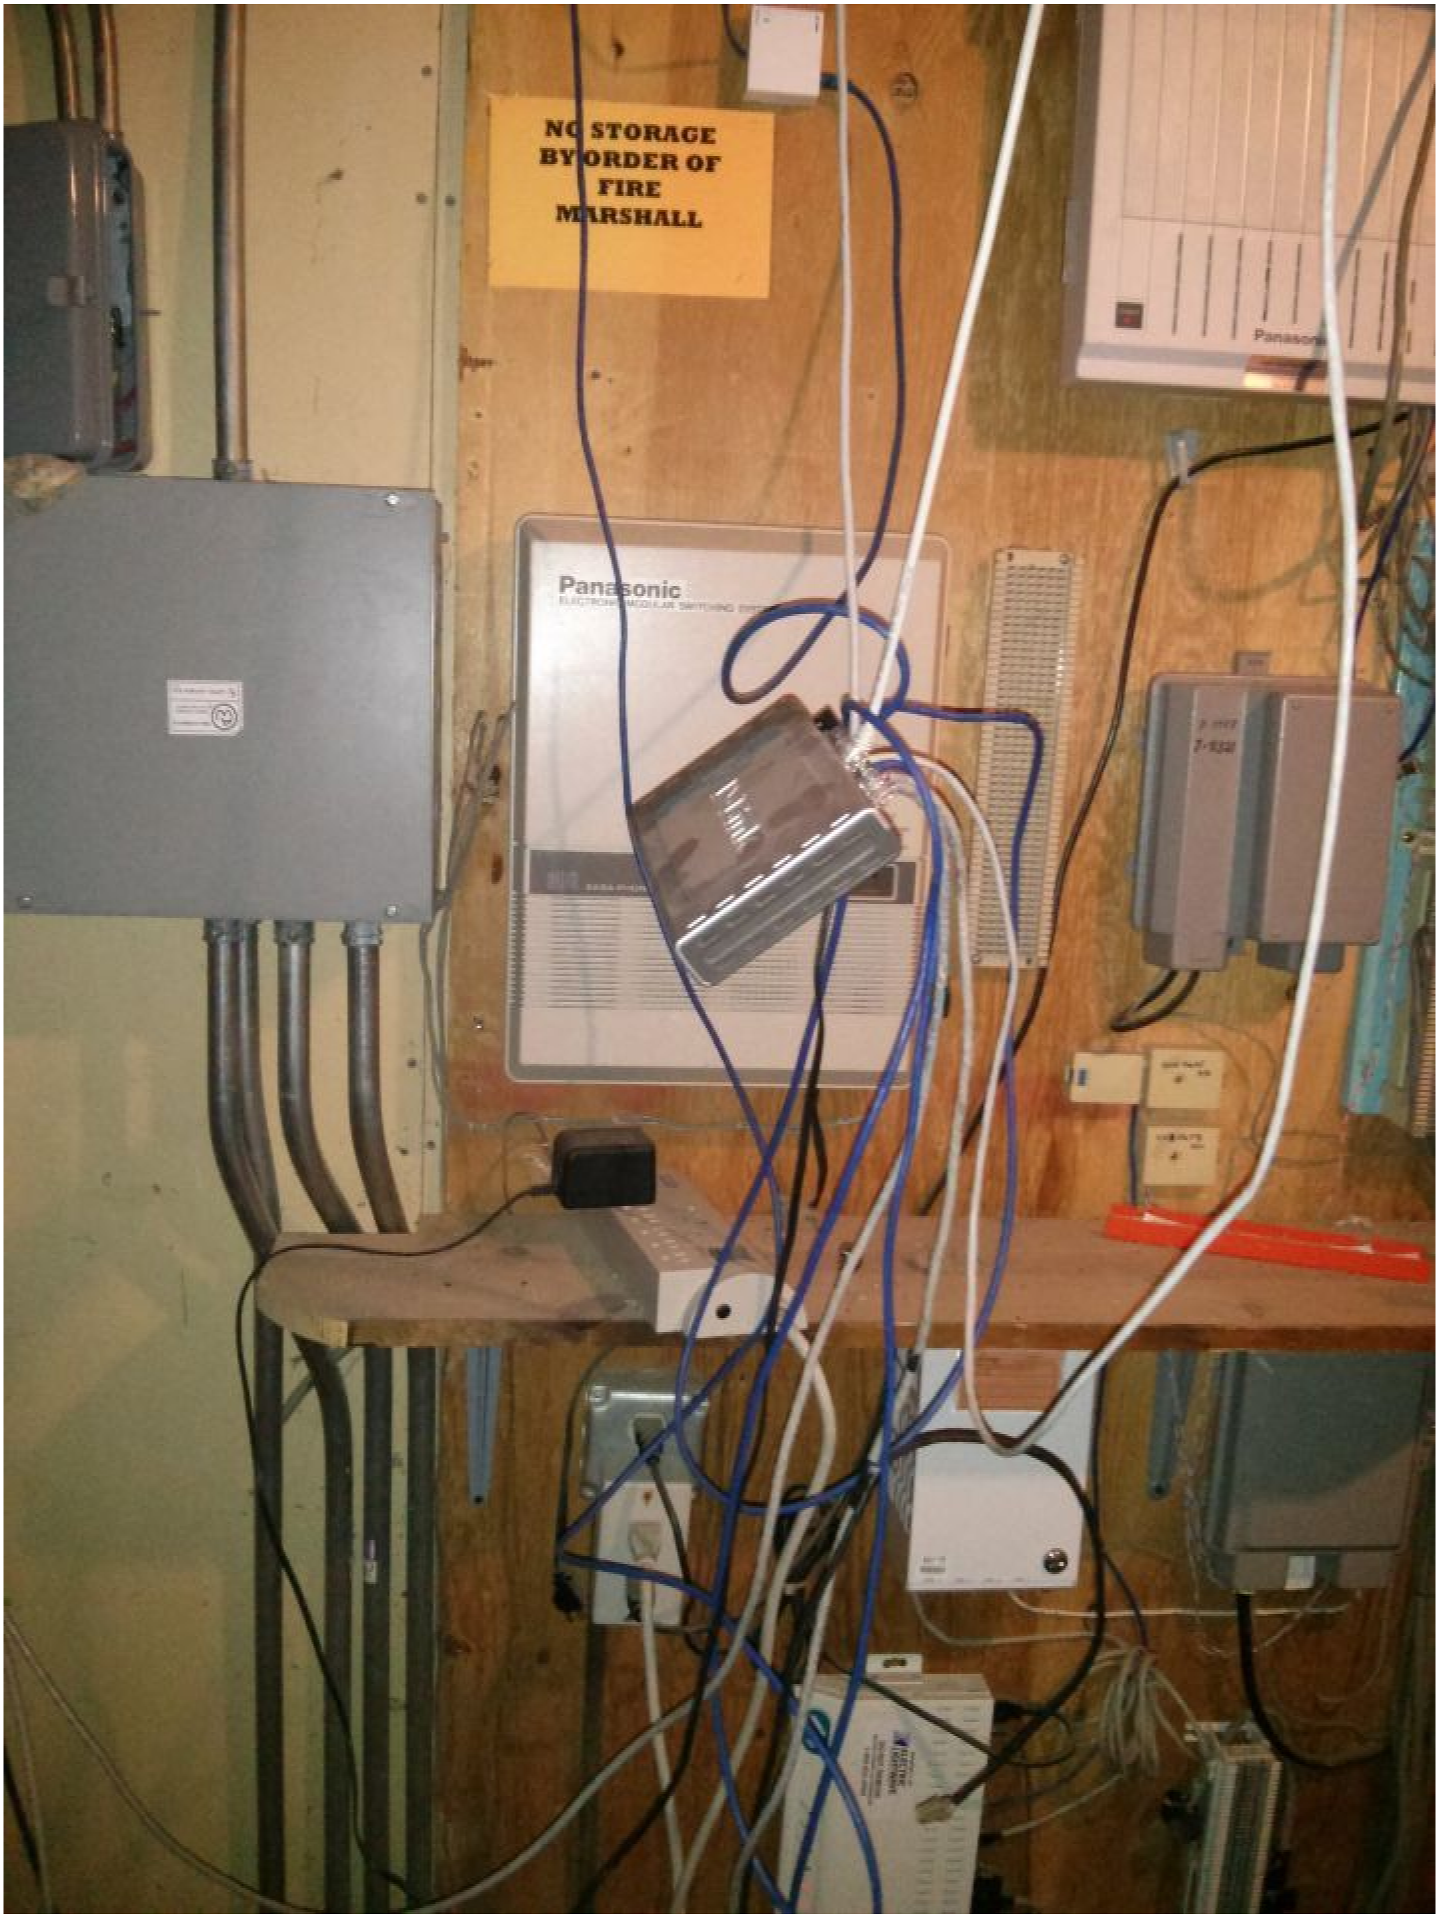
NO STORAGE
BY ORDER OF
FIRE
MARSHALL
Panasonic
ELECTRONIC MODULAR SWITCHING SYSTEM
Panasonic

## Security News

**CVE-2018-17612**
"Certificate Management Vulnerability in Sennheiser HeadSetup"
https://www.secorvo.de/publikationen/headsetup-vulnerability-report-secorvo-2018.pdf

Remember the Lenovo SuperFish debacle from 2014?

Lenovo began bundling the SuperFish Adware with some of its computers in September 2014. On February 20, 2015, the United States Department of Homeland Security advised its removal... along with its associated root certificate.

Lenovo came under fire in 2014 for pre-loading the code, which powered something called VisualDiscovery. This was meant to help shoppers by analyzing images on the web and presenting similar product offers with lower prices—thus "helping users search for images without knowing exactly what an item is called or how to describe it in a typical text-based search engine."

But to do this, the adware needed to intercept, decrypt and inspect all web browser connections, even HTTPS/TLS-protected communications. To do this, it installed its own self-signed root certificate to allow it to impersonate other websites to the PC user's browser.

That behavior would have been troubling enough. But what's worse is that rather than minting a unique and custom per-installation certificate, the same single "SuperFish" root certificate was installed onto every Lenovo PC which carried the pre-installed SuperFish adware.

Now, it's true that our machines already have a bunch of self-signed root certs. That's what trusted CA certs are. But what makes this system safe is that the server which is asserting its identity with a CA-signed cert is able to keep the certificate's associated private key safe -- since it's safely locked up at the far end of the connection.

Even those annoying TLS-intercepting "middleboxes", which enterprises use to peer into all of their traffic coming into and out of their intranet, are self-contained and located in a secured environment.

What made the SuperFish transgression so bad was that the server software was right there in the same PC since it was performing the equivalent of a local man-in-the-middle attack. And THAT meant that the certificate's matching private key HAD to also be right there in the PC.

And since the private key needed to be used for the software to operate, its own usage decryption key also needed to be present. So it wasn't very startling when Errata Security's CEO Robert Graham announced that he had been able to crack the private key for the SuperFish certificate... thus effectively breaking the HTTPS security for ALL affected Lenovo laptops.

And now, in the news, we learn that the slowly turning wheels of justice have finally clicked over and Lenovo has agreed to a settlement in a 32-state class-action lawsuit. A federal court has approved a large payout fund for Lenovo, which will be required to create a $7.3 million reservoir, set aside for settling a class action lawsuit over those surreptitious adware installations. Last week, the U.S. District Court for the Northern District of California granted preliminary approval for the settlement, which will pay out on 27 class action lawsuits that were consolidated in June 2015 into a single action. The settlement does not include attorneys' fees, so it's likely that Lenovo will see its costs edge even further upward.

So there's that.

But what mostly put all of this back on the map for this podcast is that... believe it or not... history has repeated itself. ArsTechnica's headline reads: "Sennheiser discloses monumental blunder that cripples HTTPS on PCs and Macs". BleepingComputer titled their coverage: "Sennheiser Headset Software Could Allow Man-in-the-Middle SSL Attacks"

And, yes, you heard that right. This same mistake was made not in trade for some grand sweeping advertising scheme... but only so that the Sennheiser Headset software could run in the user's web browser and securely connect to their own local software also running in the same machine. In other words, with web browsers becoming increasingly fanatical about HTTPS over HTTP -- even when the connection is to the same PC at the unroutable localhost IP of 127.0.0.1 -- Sennheiser decided that the easiest way to skin that cat would be to run their own local secure webserver on the user's machine. And just as with SuperFish before it, that meant that the matching "supposed to be kept super-secret" private key needed to be present too... thus making it available for discovery and exfiltration.

And Sennheiser committed the same grevious error as SuperFish: Rather than minting a custom per-application certificate, all installations use the same one.

The guys who discovered this -- Secorvo Security Consulting -- gen'ed up a custom certificate of their own. In their report they wrote: "Then we created a new key pair and used our fraudulent CA to issue a wildcard TLS server certificate for hosts in the domains of Google, Sennheiser and some of Sennheiser's competitors..." ... thus allowing them to remotely impersonate any Google and other domains for anyone who had ever installed that Sennheiser software.

"Ever"... yep... because as if all that wasn't bad enough, the Sennheiser root certificate was left behind, still installed in the system's root cert store after the Sennheiser software is completely uninstalled.

Since dropping certs, whose private keys are readily obtained, into the Windows root store is uncool, Microsoft has gotten into the remediation act:

ADV180029 | Inadvertently Disclosed Digital Certificates Could Allow Spoofing
https://portal.msrc.microsoft.com/en-us/security-guidance/advisory/ADV180029

Microsoft is publishing this advisory to notify customers of two inadvertently disclosed digital certificates that could be used to spoof content and to provide an update to the Certificate Trust List (CTL) to remove user-mode trust for the certificates. The disclosed root certificates were unrestricted and could be used to issue additional certificates for uses such as code signing and server authentication. More details are here: Certificate Management Vulnerability in Sennheiser HeadSetup and the CVE is here: CVE-2018-17612.

The certificates were inadvertently disclosed by the Sennheiser HeadSetup and HeadSetup Pro software. Customers who installed this software may be vulnerable, and should visit HeadSetup Update for an updated version of the HeadSetup & HeadSetup Pro software.

As a precaution, Microsoft has updated the Certificate Trust List to remove user-mode trust for these certificates. Customers who have not installed Sennheiser HeadSetup software have no action to take to be protected. Customers who have installed Sennheiser HeadSetup software should update that software via the links above.

SysInternals: SigCheck
https://docs.microsoft.com/en-us/sysinternals/downloads/sigcheck
https://download.sysinternals.com/files/Sigcheck.zip
Command:  SigCheck -tv

**Marriott and their half a BILLION (with a 'B') accounts**
https://answers.kroll.com/

“Starwood Guest Reservation Database Security Incident”

Marriott purchased the Starwood properties in 2016.
The Starwood properties include:  W Hotels, St. Regis, Sheraton Hotels & Resorts, Westin Hotels & Resorts, Element Hotels, Aloft Hotels, The Luxury Collection, Tribute Portfolio, Le Méridien Hotels & Resorts, Four Points by Sheraton and Design Hotels. Starwood branded timeshare properties are also included.

The beginning of the security incident reads:

Marriott values our guests and understands the importance of protecting personal information. We have taken measures to investigate and address a data security incident involving the Starwood guest reservation database. The investigation has determined that there was unauthorized access to the database, which contained guest information relating to reservations at Starwood properties on or before September 10, 2018. This notice explains what happened, measures we have taken, and some steps you can take in response.

On September 8, 2018, Marriott received an alert from an internal security tool regarding an attempt to access the Starwood guest reservation database. Marriott quickly engaged leading security experts to help determine what occurred. Marriott learned during the investigation that there had been unauthorized access to the Starwood network since 2014. Marriott recently discovered that an unauthorized party had copied and encrypted information, and took steps towards removing it. On November 19, 2018, Marriott was able to decrypt the information and determined that the contents were from the Starwood guest reservation database.

Marriott has not finished identifying duplicate information in the database, but believes it contains information on up to approximately 500 million guests who made a reservation at a Starwood property. For approximately 327 million of these guests, the information includes some combination of name, mailing address, phone number, email address, passport number, Starwood Preferred Guest (“SPG”) account information, date of birth, gender, arrival and departure information, reservation date, and communication preferences. For some, the information also includes payment card numbers and payment card expiration dates, but the payment card numbers were encrypted using Advanced Encryption Standard encryption (AES-128). There are two components needed to decrypt the payment card numbers, and at this point, Marriott has not been able to rule out the possibility that both were taken. For the remaining guests, the information was limited to name and sometimes other data such as mailing address, email address, or other information. Marriott reported this incident to law enforcement and continues to support their investigation. We have already begun notifying regulatory authorities.

Marriott deeply regrets this incident happened. From the start, we moved quickly to contain the incident and conduct a thorough investigation with the assistance of leading security experts. Marriott is working hard to ensure our guests have answers to questions about their personal information with a dedicated website and call center. We are supporting the efforts of law enforcement and working with leading security experts to improve. Marriott is also devoting the resources necessary to phase out Starwood systems and accelerate the ongoing security enhancements to our network.

Marriott has established a dedicated call center to answer questions you may have about this incident. The call center is open seven days a week and is available in multiple languages. Our dedicated call center may experience high call volume initially, and we appreciate your patience.

The Marriott breach is big, but it's still only #2, behind Yahoo's massive 3 Billion user breach back in 2016.

**Senators call for data security law in wake of Marriott breach**
https://www.cnet.com/news/lawmakers-call-for-data-security-legislation-in-wake-of-marriott-breach

They want federal legislation to protect consumers and hold companies accountable.

Senator Mark Warner said in a statement: "We must pass laws that require data minimization, ensuring companies do not keep sensitive data that they no longer need. And it is past time we enact data security laws that ensure companies account for security costs rather than making their consumers shoulder the burden and harms resulting from these lapses."

The "data minimization" makes a lot of sense to me. It would be real protection and unlike other policies, its enforcement is it is easily tested: When a breach occurs and the information that has escaped is revealed, the responsible company can be asked: "Why exactly were you still retaining all of that data from a decade ago?  What was its business purpose?"

**Meanwhile, a hack of the Internet's leading Question & Answer site, Quora, was just announced:**
https://blog.quora.com/Quora-Security-Update

Adam D'Angelo posting on the Quora blog with the heading: "Quora Security Update"
We recently discovered that some user data was compromised as a result of unauthorized access to one of our systems by a malicious third party. We are working rapidly to investigate the situation further and take the appropriate steps to prevent such incidents in the future.

We also want to be as transparent as possible without compromising our security systems or the steps we're taking, and in this post we'll share what happened, what information was involved, what we're doing, and what you can do.

We're very sorry for any concern or inconvenience this may cause.

What happened

On Friday we discovered that some user data was compromised by a third party who gained unauthorized access to one of our systems. We're still investigating the precise causes and in addition to the work being conducted by our internal security teams, we have retained a leading digital forensics and security firm to assist us. We have also notified law enforcement officials.

While the investigation is still ongoing, we have already taken steps to contain the incident, and our efforts to protect our users and prevent this type of incident from happening in the future are our top priority as a company.

What information was involved

For approximately 100 million Quora users, the following information may have been compromised:

- Account information, e.g. name, email address, encrypted (hashed) password, data imported from linked networks when authorized by users
- Public content and actions, e.g. questions, answers, comments, upvotes
- Non-public content and actions, e.g. answer requests, downvotes, direct messages (note that a low percentage of Quora users have sent or received such messages)

Questions and answers that were written anonymously are not affected by this breach as we do not store the identities of people who post anonymous content.

The overwhelming majority of the content accessed was already public on Quora, but the compromise of account and other private information is serious.

What we are doing

While our investigation continues, we're taking additional steps to improve our security:

- We’re in the process of notifying users whose data has been compromised.
- Out of an abundance of caution, we are logging out all Quora users who may have been affected, and, if they use a password as their authentication method, we are invalidating their passwords.
- We believe we’ve identified the root cause and taken steps to address the issue, although our investigation is ongoing and we’ll continue to make security improvements.

We will continue to work both internally and with our outside experts to gain a full understanding of what happened and take any further action as needed.

What you can do

We’ve included more detailed information about more specific questions you may have in our help center, which you can find here:
https://help.quora.com/hc/en-us/articles/360020212652

If you were affected, we will update you with relevant details via email.

While the passwords were encrypted (hashed with a salt that varies for each user), it is generally a best practice not to reuse the same password across multiple services, and we recommend that people change their passwords if they are doing so.

Conclusion

It is our responsibility to make sure things like this don’t happen, and we failed to meet that responsibility. We recognize that in order to maintain user trust, we need to work very hard to make sure this does not happen again. There’s little hope of sharing and growing the world’s knowledge if those doing so cannot feel safe and secure, and cannot trust that their information will remain private. We are continuing to work very hard to remedy the situation, and we hope over time to prove that we are worthy of your trust.

**And in other news... the records of 114 Million US Citizen and Companies Exposed Online**
https://blog.hackenproof.com/industry-news/new-data-breach-exposes-57-million-records

New Data Breach exposes 57 million records

A massive 73 GB data breach was discovered during a regular security audit of publicly available servers with the Shodan search engine. Prior to this publication, there were at least 3 IPs with the identical Elasticsearch clusters misconfigured for public access. First IP was indexed by Shodan on November 14th, 2018. An open Elasticsearch instance exposed personal info of 56,934,021 US citizens, with information such as first name, last name, employers, job title, email, address, state, zip, phone number, and IP address.

Another index of the same database contained more than 25 million records with more of a “Yellow Pages” details directory: name, company details, zip address, carrier route, latitude/longitude, census tract, phone number, web address, email, employees count, revenue numbers, NAICS codes, SIC codes, and etc.

While the source of the leak was not immediately identifiable, the structure of the field ‘source’ in data fields is similar to those used by a data management company Data & Leads Inc. However, we weren’t able to get in touch with their representatives.

Moreover, shortly before this publication Data & Leads website went offline and now is unavailable.

As of today, the database is no longer exposed to the public, however, it is unknown for how long it has been online before Shodan crawlers indexed it on November 14th and who else might have accessed the data.

Importance of Responsible Disclosure

Our goal is to help protect data on the Internet by identifying data leaks and following responsible disclosure policies. Our mission is to make the cyber world safer by educating businesses and communities worldwide on ethical vulnerability disclosure policy (VDP). We regularly publish reports on data leak discoveries made by our research team, you can find them on our blog.

We have previously reported that the lack of authentication allowed the installation of malware or ransomware on the Elasticsearch servers. The public configuration allows the possibility of cybercriminals to manage the whole system with full administrative privileges. Once the malware is in place criminals could remotely access the server resources and even launch a code execution to steal or completely destroy any saved data the server contains.

According to the DB-Engines ranking, Elasticsearch is the most popular enterprise search engine.

**What does all of this hacking mean?**

As usual it's not monolithic.  There's a range of vulnerability and responsibility and accountability.

It sounds as though these "Data & Leads" guys are a data broker who were siphoning up consumer information (the Data) for resale (the Leads) and they were being irresponsible custodians of this data that wasn't really there's to have or hold in the first place.

It sounds as though Quora's system was designed correctly and responsibly, but that bad guys found a way inside, even so. They appear to have responded to the incident quickly and responsibly. It would be nice to know what the breach point is -- and we know they are desperate to know, if they don't already. But with what information we have, they appear to have and be acting honorably.

And it sounds as though Marriott made a somewhat diseased property acquisition two years ago which came back to bite them. In retrospect, they should have performed a careful forensic analysis of their Starwood acquisition beforehand. And who knows, perhaps they did and it was missed? If it was just an analysis for show, then perhaps they should have spent more to make it real. But we also know from the Sony incident that it's possible for an APT -- an advanced persisten threat -- actor to lodge themselves into a large organization and to readily remain hidden for years.

I think that the best we can do, proactively, as consumers is be aware of the constant possibility that these details of our lives, which we would like to have remain private, are increasingly unlikely to remain so -- at least in part.

I believe that, where feasible, everyone should keep their credit bureau reporting data locked from access so that identity thieves are unable to successfully apply for and abuse credit in our names.

And large sums of cash should not be left lying around in electronically accessible piles in a bank account. Build a firewall between any money being stored and money being actively used for payments. Banks just shrug when a bad guy electronically transfers your funds out of your accounts.

We already know that passwords should not be reused, and that individually unique passwords must have a large amount of entropy. So this means that a password manager is no longer an option.

And as most of us now do, it's useful to separate our online lives into trust silos so that rather than presenting one large monolithic target to the online world, we appear to be many separate virtual individuals with differing names and eMail addresses and other data. That way, the breach of any one of them is not a breach of our entire online identity.

It's truly unfortunate that all of that is recommended. But this is the world we live in today.

**EternalSilence: "EternalBlue" + "Silent Cookie"**
Akamai says that over 45,000 routers have been compromised already.

Akamai has detected an ingenious malware campaign that alters configurations on home and small office routers to open connections INWARD toward internal networks so crooks can access and infect previously isolated computers.

Hackers achieve this via the UPnProxy technique we've talked about before. UPnProxy exploits vulnerabilities in the UPnP services installed on some routers to alter the device's NAT (Network Address Translation) tables.

Last April, we saw that hackers were using this technique to convert routers into proxies for bouncing DDoS, spamming, phishing and other traffic. But in a report published last Wednesday, Akamai says it's seen a new variation of UPnProxy where hackers are now leveraging UPnP services to insert different rules into routers NAT tables: These rules still function as proxy redirectors, but instead of relaying traffic externally, they allow an external hacker to connect to the SMB ports (139, 445) of devices and computers located behind the router, on the internal network.

Akamai says that from the 277,000 routers with vulnerable UPnP services exposed online, 45,113 have been modified in the recent campaign they have uncovered. They found that one particular hacker, or hacker group, has spent weeks creating a custom NAT entry named 'galleta silenciosa' ('silent cookie/cracker' in Spanish) on these 45,000 routers.

Akamai says it detected "millions of successful injections" during which crooks connected through these ports to devices beyond the routers. Akamai put the number of these devices around the 1.7 million figure.

What the hackers did, Akamai can't tell, as they don't have visibility inside those networks. But the company is quite certain these "injections" have something to do with EternalBlue, one of the pieces of malware developed by the US National Security Agency, and which leaked online last year, and the malware that was at the heart of the WannaCry and NotPetya ransomware outbreaks.

Furthermore, Akamai also believes hackers deployed EternalRed, a variant of EternalBlue that can infect Linux systems via Samba, the SMB protocol implementation for Linux.

There is some kinda good news, however, since this doesn't appear to be a nation-state orchestrated hacking operation with a bigger end goal in mind. Akamai said that "Recent scans suggest that these attackers are opportunistic intruders. The goal isn't a targeted attack. It's an attempt at leveraging off the shelf exploits to cast a wide net into a relatively small pond in the hope of scooping up a pool of previously inaccessible devices."

Companies and individuals who don't want to be victims of these and future attacks are advised to either disable the UPnP service on their routers -- PLEASE PLEASE!! -- or if UPnP must be exposed on the WAN side, at least obtain a well-secured router that doesn't use a vulnerable UPnP implementation.

Akamai refers to this particular router hacking campaign as EternalSilence, a name derived from the use of the EternalBlue exploits and Silent Cookie, the name of the malicious NAT table entries. The company has also published instructions at the bottom of its report on how to remove the malicious NAT table entries from affected routers.

**So:  You're apparently a =BIG= fan if the Swedish YouTuber, comedian and video game commentator Felix Kjellberg, known on YouTube as PewDiePie.**

You learn that his #1 YouTube position by subscriber count is being endangered by "T-Series", some lame Indian music record label and film company that simply uploads videos of Bollywood trailers and songs. Both YouTubers’ channels have more than 73 million subscribers, but at the moment PewDiePie is currently leading by a comparatively narrow margin of just 300,000.

You decide to take matters into your own hands!

You want to send out a message... to everyone;  But you don't have President Trump's "the world is ending" universal cellular phone Presidential SMS blaster code. But you ARE in possession of some modest hacking skillz; and...

You're not very troubled by subtle questions of ethics, morality or legality.

How do you proceed?

You surf on over to Shodan and poke around to find a bunch of something.

You find a lot of printers. Publicly exposed printers. Perfect. PewDiePie, here we come!

According to the hacker, he found three different vulnerable printing protocols on Shodan (IPP, LPD, and JetDirect) with up to 800,000 vulnerable printers in total.

He tweeted: “I was horrified to see over 800,000 results show up in total. I was baffled, but determined to try this. So I picked the first 50,000 printers I found running on port 9100 and downloaded the list off Shodan.”

The hacker then used PRET -- the PRinter Exploitation Toolkit -- on Github – which gives hackers the ability to access files, damage the printer, or access the internal network.

However, @HackerGiraffe said that he only wanted to use the kit to print out messages about PewDiePie... to spread awareness.

He tweeted: “PRET [Printer Exploitation Toolkit] had the scariest of features. Ability to access files, damage the printer, access the internal network…things that could really cause damage. So I had to do this, to at least help organizations and people that can protect themselves,” he said in a Tweet.

Ah... so it doubled as a public service announcement.

The hacker typed up a bash script, which runs an exploit kit against the impacted IP with commands to print a message then quit. He then uploaded the script onto his server and left it running.

The printed message said: “PewDiePie is in trouble and he needs your help to defeat T-Series! PewDiePie, the currently most subscribed to channel on Youtube, is at stake of losing his position as the number one position by an Indian company called T-Series, that simply uploads videos of Bollywood trailers and songs.”

The message then urged readers to unsubscribe from T-Series and subscribe to PewDiePie, and concluded the message by telling readers to tell everyone they know….

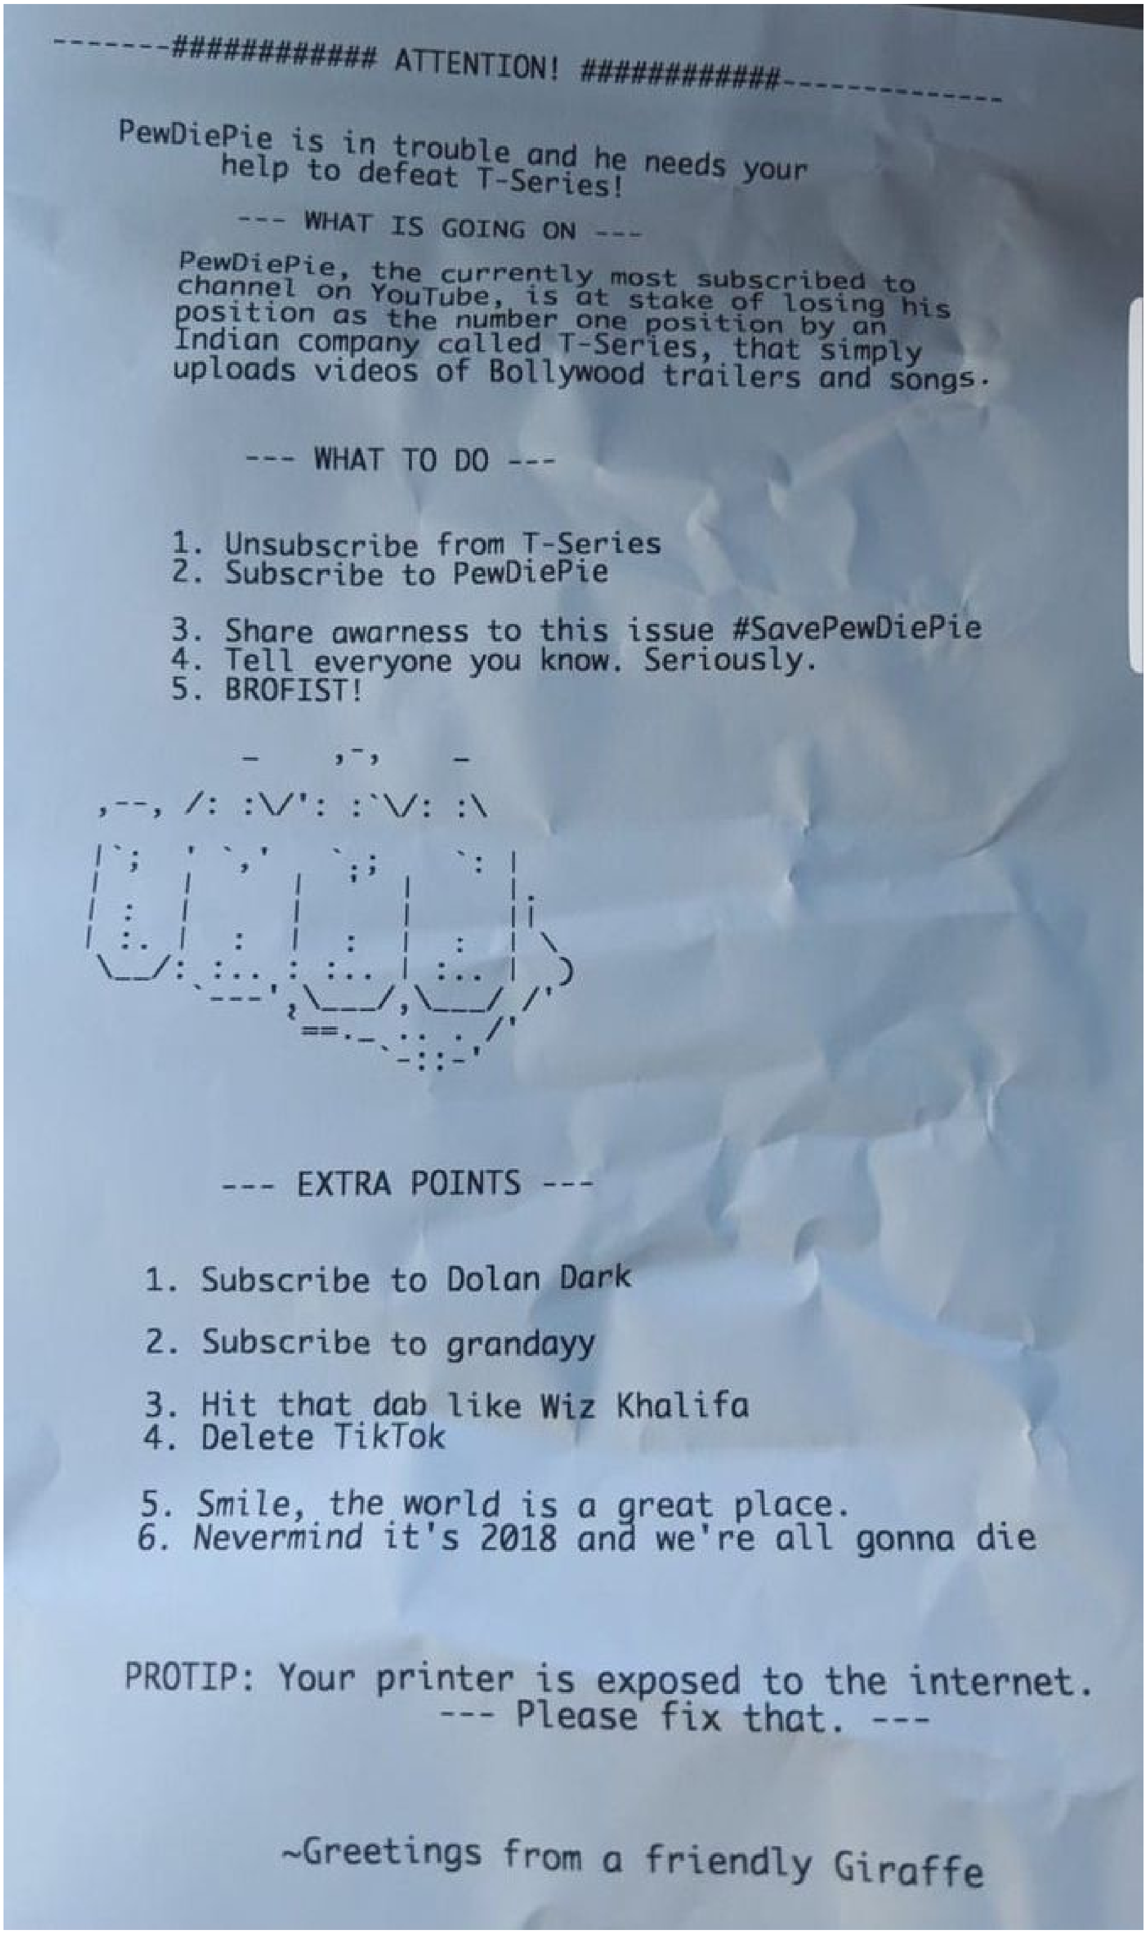

-------############ ATTENTION! ############-------------

PewDiePie is in trouble and he needs your help to defeat T-Series!

--- WHAT IS GOING ON ---

PewDiePie, the currently most subscribed to channel on YouTube, is at stake of losing his position as the number one position by an Indian company called T-Series, that simply uploads videos of Bollywood trailers and songs.

--- WHAT TO DO ---

1. Unsubscribe from T-Series
2. Subscribe to PewDiePie

3. Share awarness to this issue #SavePewDiePie
4. Tell everyone you know. Seriously.
5. BROFIST!

```
        _     ,-,     _
 ,--, /: :\/': :`\/: :\
|`;   ' `,'   `.;   `: |
|    |     |  '  |     |.
| :  |     |     |     ||
| :. |  :  |  :  |  :  | \
 \__/: :.. : :.. | :.. |  )
      `---',\___/,\___/ /'
           `==._ .. . /'
                `-::-'
```

--- EXTRA POINTS ---

1. Subscribe to Dolan Dark

2. Subscribe to grandayy

3. Hit that dab like Wiz Khalifa
4. Delete TikTok

5. Smile, the world is a great place.
6. Nevermind it's 2018 and we're all gonna die

PROTIP: Your printer is exposed to the internet.
--- Please fix that. ---

~Greetings from a friendly Giraffe

**And the PewDiePie hack has apparently spawned a new web service over the weekend: Printer-Spam-as-a-Service.**
https://www.zdnet.com/article/new-online-service-will-hack-printers-to-spew-out-spam/

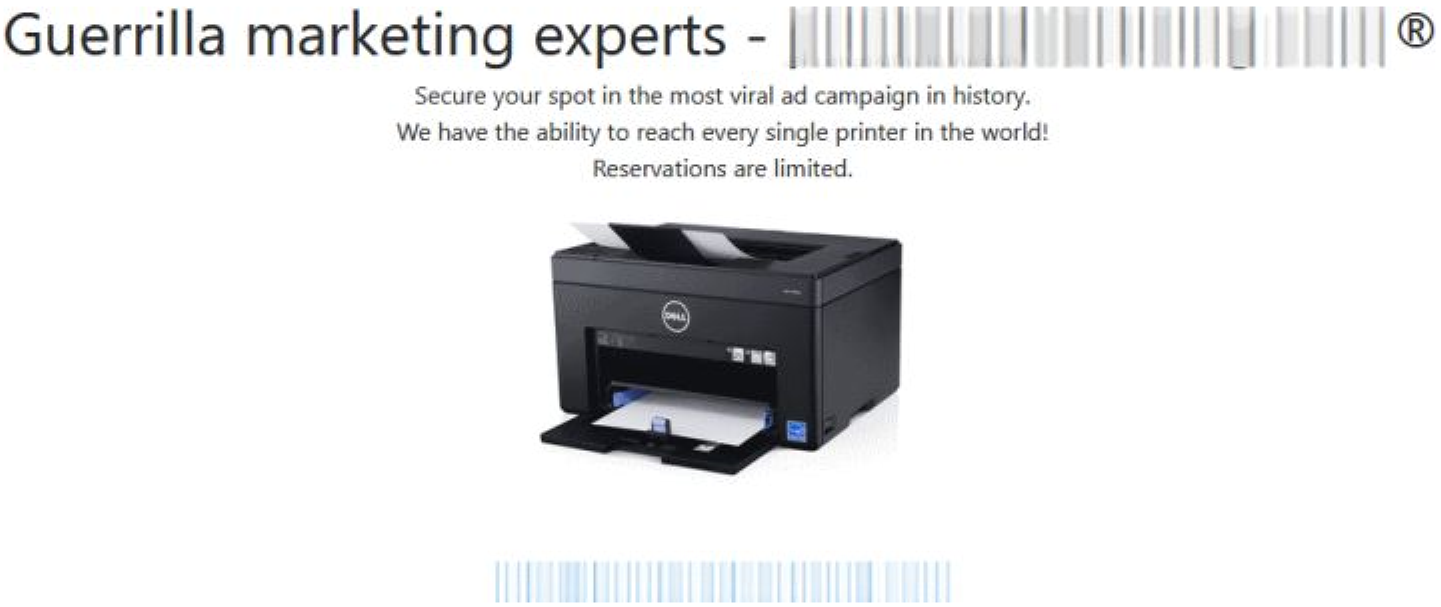

Guerrilla marketing experts - [redacted] ®

Secure your spot in the most viral ad campaign in history.
We have the ability to reach every single printer in the world!
Reservations are limited.

[redacted]

"Printer Advertising" is the not-very-imaginative name given to an upstart service which is proposing to mass-push printed spam messages to exposed Internet printers, for a price.

The good news is, the end result will be the removal of at least some of Shoran's inventory of up to 800,000 currently exposed, accessible and perhaps vulnerable printers from the public Internet.

Andrew Morris, the founder of GreyNoise Intelligence, detected the message in one of his company's honeypots on Sunday, but the spam campaign pushing this ad to Internet-connected printers has continued through yesterday.

The printre spam all originates from an IP address which is quite well known to those who monitor this sort of "Internet Background Radiation" (the term I coined for this long ago). The IP address is: 194.36.173.50 ... which is known for generating quite a lot of bad traffic. That IP is continuously scanning for router UPnP services, ColdFusion plugins, LDAP, web, DNS, and Memcached servers.

Making a Ransomware Payment? It May Now Violate U.S. Sanctions
https://www.bleepingcomputer.com/news/security/making-a-ransomware-payment-it-may-now-violate-us-sanctions/

Thinking about making a ransomware payment? If so, you may want to think twice before doing so as it could land you in trouble for violating U.S. government sanctions.

This week the Department of Justice unsealed a grand jury indictment against two Iranian hackers allegedly responsible for the SamSam Ransomware. As part of this indictment, for the first time the U.S. Department of the Treasury’s Office of Foreign Assets Control (OFAC) also publicly attributed cryptocurrency addresses to individuals who were involved in converting ransomware cryptocurrency payments to fiat currency.

The Department of Treasury's announcement stated: "While OFAC routinely provides identifiers for designated persons, today’s action marks the first time OFAC is publicly attributing digital currency addresses to designated individuals"

In this particular case, the cryptocurrency addresses are being attributed to Iran-based individuals named Ali Khorashadizadeh and Mohammad Ghorbaniyan who the U.S. government states have facilitated the exchange of ransomware payments into Iranian Rial.

The addresses attributed to these individuals are:  1AjZPMsnmpdK2Rv9KQNfMurTXinscVro9V and 149w62rY42aZBox8fGcmqNsXUzSStKeq8C and contain a combined total of 5,901 bitcoins. Which puts the value of that cache at over $23 million USD.

OFAC has also added Khorashadizadeh and Ghorbaniyan to the Specially Designated Nationals And Blocked Persons List (SDN), which means that U.S. individuals and companies are blocked from doing business or conducting any transactions with these individuals. These sanctions could also affect non U.S. businesses and individuals who conduct transactions with them due to secondary sanctions.

**iOS Fitness Apps Robbing Money From Apple Victims**
https://www.welivesecurity.com/2018/12/03/scam-ios-apps-promise-fitness-steal-money-instead/

iOS apps used Touch ID feature to trick users into paying hefty fees

Be a bit wary of apps that ask for your fingerprint to access their own stored data. Two apps have been found in (and now removed from) the iOS App Store using that fingerprint to authorize an Apple Pay transaction of $99 to $119.

The two apps were “Fitness Balance” and “Calories Tracker" and they would each state that their users needed to supply their fingerprint to access the app's features.

The apps had multiple 5-star reviews and had established fake good looking reputations. So users could certainly be fooled into believing and trusting them.

ESET security explained in their posting:

After a user fires up any of the abovementioned apps for the first time, the apps request a fingerprint scan to “view their personalized calorie tracker and diet recommendations”. Only moments after the user complies with the request and places their finger on the fingerprint scanner, the apps then display a pop-up showing a dodgy payment amounting to 99.99, 119.99 USD or 139.99 EUR.

This pop-up is only visible for about a second, however, if the user has a credit or debit card directly connected to their Apple account, the transaction is considered verified and money is wired to the operator behind these scams.

iPhone X users would be protected if they enabled "Double Click Side Button" (which was enabled by default on my iPhone) since the 'X' has no home button. Earlier iPhone model users who do have a home button would be charged automatically. if they had enough credit or a saved credit card and Touch ID was enabled.

The good news is that credit card users are indemnified against such fraudulent charges, but it's likely that the bad guys made some money before they were shutdown.

**Mozilla is expanding its DoH --> DNS over HTTPS**
https://blog.mozilla.org/futurereleases/2018/11/27/next-steps-in-dns-over-https-testing/

"Next Steps in DNS-over-HTTPS Testing"

Over the past few months, Mozilla has experimented with DNS-over-HTTPS (DoH). The intention is to fix a part of a DNS ecosystem that simply isn't up to the modern, secure standards that every Internet user should expect. Today, we want to let you know about our next test of the feature.

Our initial tests of DoH studied the time it takes to get a response from Cloudflare's DoH resolver. The results were very positive – the slowest users show a huge performance improvement. A recent test in our Beta channel confirmed that DoH is fast and isn't causing problems for our users. However, those tests only measure the DNS operation itself, which isn't the whole story.

Content Delivery Networks (CDNs) provide localized DNS responses depending on where you are in the network, with the goal being to send you to a host which is near you on the network and therefore will give you the best performance. However, because of the way Cloudflare resolves names, this process works less well when you are using DoH with Firefox.

The result is that the user might get less well-localized results that could result in a slow user experience even if the resolver itself is accurate and fast.

This is something we can test. We are going to study the total time it takes to get a response from the resolver and fetch a web page. To do that, we're working with Akamai to help us understand more about the performance impact. Firefox users enrolled in the study will automatically fetch data once a day from four test web pages hosted by Akamai, collect information about how long it took to look up DNS and then send that performance information to Firefox engineers for analysis. These test pages aren't ones that the user would automatically retrieve and just contain dummy content.

A soft rollout to a small portion of users in our Release channel in the United States will begin this week and end next week. As before, this study will use Cloudflare's DNS-over-HTTPS service and will continue to provide in-browser notifications about the experiment so that everyone is fully informed and has a chance to decline participation in this particular experiment. Moving forward, we are working to build a larger ecosystem of trusted DoH providers, and we hope to be able to experiment with other providers soon.

We don’t yet have a date for the full release of this feature. We will give you a readout of the result of this test and will let you know our future plans at that time. So stay tuned.

**Microsoft is rumored to be abandoning it investment in its Edge Browser for Windows 10… And switching to a new browser built on Chrome!**

Zac Bowden over at Windows Central reports...

Microsoft's Edge web browser has seen little success since its debut on Windows 10 in 2015. Built from the ground up with a new rendering engine known as EdgeHTML, Microsoft Edge was designed to be fast, lightweight, and secure, but it launched with a plethora of issues that resulted in users rejecting it early on. Edge has since struggled to gain traction, thanks to its continued instability and lack of mindshare, from users and web developers.

Because of this, I'm told that Microsoft is throwing in the towel with EdgeHTML and is instead building a new web browser powered by Chromium, which uses a similar rendering engine first popularized by Google's Chrome browser known as Blink. Codenamed "Anaheim," this new browser for Windows 10 will replace Edge as the default browser on the platform, according to my sources, who wish to remain anonymous. It's unknown at this time if Anaheim will use the Edge brand or a new brand, or if the user interface (UI) between Edge and Anaheim is different. One thing is for sure, however; EdgeHTML in Windows 10's default browser is dead.

Many will be happy to hear that Microsoft is finally adopting a different rendering engine for the default web browser on Windows 10. Using Chromium means websites should behave just like they do on Google Chrome in Microsoft's new Anaheim browser, meaning users shouldn't suffer from the same instability and performance issues found in Edge today. This is the first step towards revitalizing Windows 10's built-in web browser for users across PCs and phones. Edge on iOS and Android already uses rendering engines native to those platforms, so not much will be changing on that front.

In addition, Microsoft engineers were recently spotted committing code to the Chromium project to help get Google Chrome running on ARM. Perhaps some of that work will translate over to getting Anaheim running on Windows 10 on ARM, too.

I expect we'll see Microsoft introduce Anaheim throughout the 19H1 development cycle, which Insiders are currently testing in the Fast ring. This is a big deal for Windows. Microsoft's web browser should finally be able to compete alongside Chrome, Opera and Firefox, and those who are all-in with the Microsoft ecosystem will finally be getting a browser from Microsoft that works well when browsing the web.

There's still lots we don't know about Anaheim, and I'm sure we'll hear more about it officially from Microsoft in the coming weeks.

**Bing was generating a false-positive warning about VLC Player**
Bing was Warning that the VLC Media Player Site is Unsafe

@VideoLAN tweeted...

Supposedly, @bing now consider vlc-3.0.4-win64.exe as a malware, which gives an annoying popup. This appeared 2 days ago, and we have no clue how to fix it (yet). We've checked, and the binary has not changed and is still correctly signed….

> tbc…
>
> — VideoLAN (@videolan) November 27, 2018

## Miscellany

**Ian Wills @ian_wills:**
GDPR (Greatly Disproportionate Privacy Response)

**Humble Bundle Cyber Security Packt Books**
https://www.humblebundle.com/books/cybersecurity-packt-books
~Just under six days remaining
DRM-free / Multi-format / PDF, ePub, Mobi

- NMAP: Network Exploration and Security Auditing
- Network Analysis using Wireshark 2
- Cryptography in Python
- Hands-On Penetration testing on Windows
- Metasploit Penetration Testing Cookbook
- Mastering pfSense
- Mastering Kali Linux
- Metasploit for Beginners
- Mastering Linux Security and Hardening

## SpinRite

Anthony May, is an E.E. from .au in SF.
https://www.quora.com/What-needs-repair-on-a-computer-that-is-harder-than-you-think

Hard Drives and their ‘bad sectors’.

Yes, ‘spinning rust’ hard-drives whose design goes back to the 1950s, still going strong throughout this decade, despite the rise of Solid State Drives, and will continue well into the next decade.

To be clear, there’s not much real “repair” goes on with modern computers, or tech in general these days, it’s deemed ‘not economically viable to repair’, compared to the cost of replacement, swapping out some module/card/motherboard/drive/etc.

But hard-drives? They’re spinning death frisbees just waiting to gobble your precious data stored on them, and more goes wrong in them than most people realise, their data kept “safe” only thanks to mathematics.

“Bad Sectors” is a term you hear occasionally, but it’s rare to hear someone who actually knows how to ‘fix’ them, because in reality there’s only one way I know that has any likelihood of fixing the problem (unless you opt for thousands of dollars at a specialty data recovery business), and it’s a commercial software utility, which automatically makes people dubious - “How could software fix hardware?!” they scoff.

...

Remember, a hard-drive doesn’t know there’s a problem with a sector of data until it tries to read it and it discovers that the math doesn’t add up any more.

...

Again respecting the analog-y nature of the magnetic alignment of ferrous particles on the surface of the hard-drive and that they can weaken over time, you can ‘exercise’ the physical hard-drive medium by reading the data from the sector, inverting its 1s & 0s and re-writing that back to the sector, then reading that back, re-inverting its 1s & 0s and finally writing that back to the sector - the net result is the data is exactly the same, but you’ve pushed every bit through a write of a 1 and then a 0 then a 1, or vice-versa, leaving the sector freshly written, but at each step the hard-drive is monitoring the error-detection math to see if there’s any sign of surface defect.

This is what SpinRite does, the commercial software I mentioned earlier. No, I’m not affiliated with Home of Gibson Research Corporation and Steve Gibson’s SpinRite software, but I’ve used it for nearly 30 years, as have countless other computer pros, and it’s rescued countless amounts of data from presumed death, because of this analog-y nature of spinning-disc hard-drives, and incredibly, the same results are now being achieved for SSDs, even though their failure mechanisms are entirely different, but the forward-error-correction math is still there, the spare sectors are still there, and so corrupted data can still be recovered, from both ‘spinning rust’ hard-drives and modern Solid State Drives.

## Closing The Loop

**Frode Burdal Klevstul @klevstul**
@SGgrc is there anything in the sqrl protocol that makes captchas obsolete?

**themainapp @themainapp**
@SGgrc @GibsonResearch just finished listening to password immortal podcast. Enjoyed your rebuttal of the paper but I think you missed a few of the papers points. I think SQRL will run into useabiltiy issues because it’s too easy.

# Stealing Webpages Rendered on Your Browser by Exploiting GPU Vulnerabilities

https://www.cc.gatech.edu/~slee3036/papers/lee:gpu.pdf

***A NEW side-channel: Residual web browser page rendering textures***

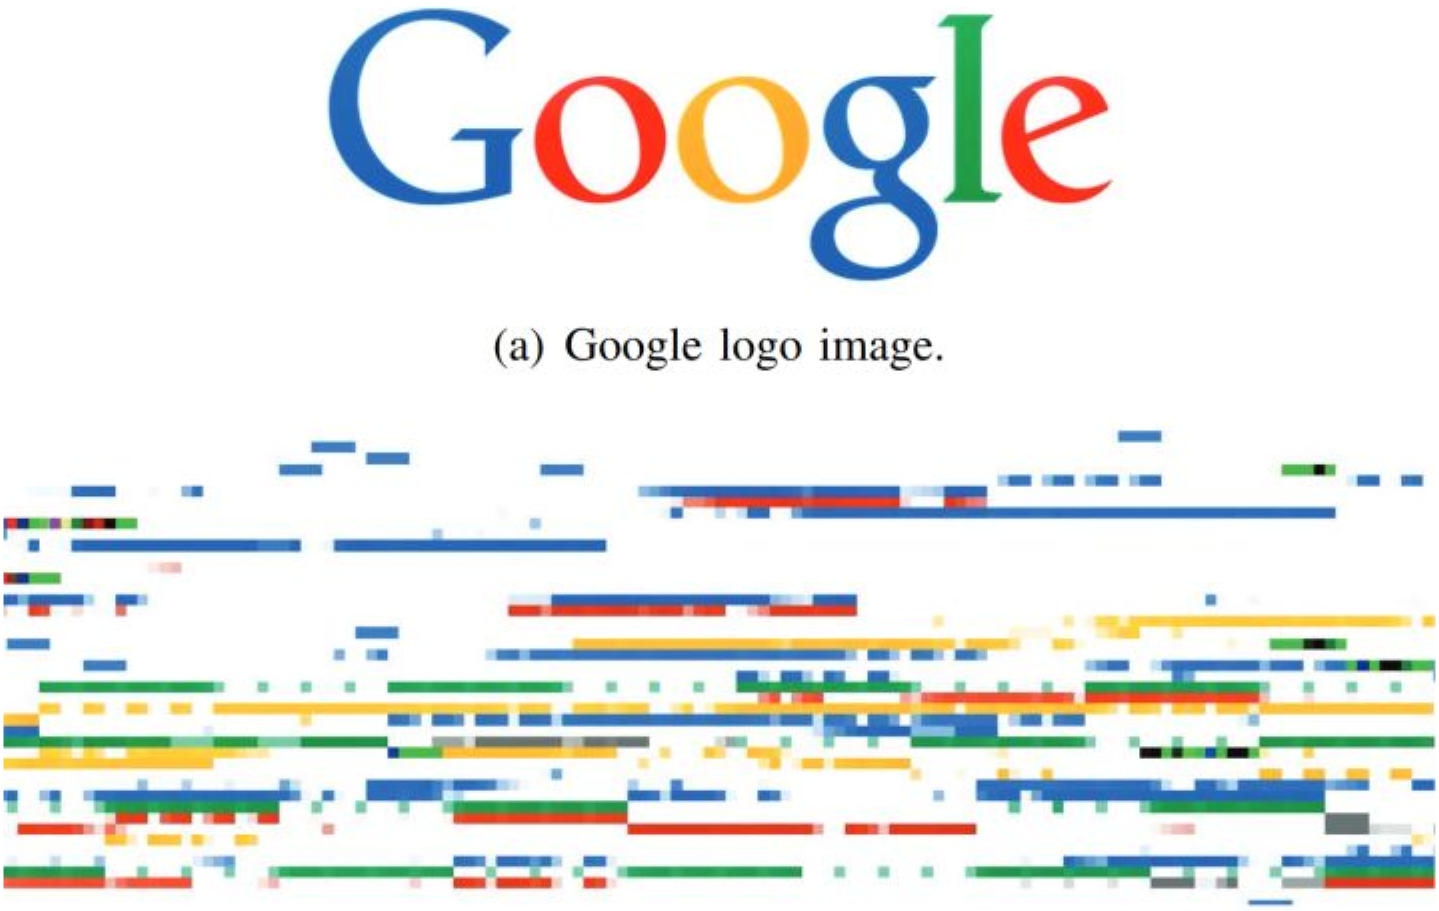

(a) Google logo image.

(b) Partial dump of Google webpage textures.

**Abstract**

Graphics processing units (GPUs) are important components of modern computing devices for not only graphics rendering, but also efficient parallel computations. However, their security problems are ignored despite their importance and popularity. In this paper, we first perform an in-depth security analysis on GPUs to detect security vulnerabilities. We observe that contemporary, widely-used GPUs, both NVIDIA's and AMD's, do not initialize newly allocated GPU memory pages which may contain sensitive user data. By exploiting such vulnerabilities, we propose attack methods for revealing a victim program's data kept in GPU memory both during its execution and right after its termination. We further show the high applicability of the proposed attacks by applying them to the Chromium and Firefox web browsers which use GPUs for accelerating webpage rendering. We detect that both browsers leave rendered webpage textures in GPU memory, so that we can infer which webpages a victim user has visited by analyzing the remaining textures. The accuracy of our advanced inference attack that uses both pixel sequence matching and RGB histogram matching is up to 95.4%.

## INTRODUCTION

This work considers how attackers can disclose sensitive data kept in graphics processing unit (GPU) memory. We aim to obtain rendered webpage textures to uncover webpages a victim user has visited. We successfully reveal such data from modern GPUs (e.g., NVIDIA and AMD GPUs) when we enable GPU-accelerated webpage rendering of recent web browsers (e.g., Chromium and Firefox). For example, Figure 1 shows the Google logo image of http://google.com and a partial dump of rendered webpage textures extracted from an NVIDIA GPU used by the Chromium web browser. Although the GPU has rearranged the textures according to its undocumented hardware characteristics, we can infer that the dump originates from the webpage because their color patterns are similar. Especially, our combined matching attack can successfully infer up to 95.4% of randomly visited 100 front pages of Alexa Top 1000 websites when a victim uses the Chromium web browser with an NVIDIA GPU.

GPUs are important and powerful components of contemporary computing devices. Personal computing devices, including desktops, laptops, and smartphones, use GPUs for supporting various graphics applications, e.g., graphical user interface (GUI), multimedia players, and video games. Large-scale computing devices, including workstations, servers, and clusters, also use GPUs for energy-efficient, massive parallel computations. GPUs utilize a large number of processing cores and a large amount of independent memory for efficiently processing graphics operations and computational workloads. For example, an NVIDIA Kepler GPU can have up to 2880 cores and 6 GB of memory, and its floating-point operation performance is nine times better than that of the recent CPUs.

Programmers can use two types of application programming interfaces (APIs) to access GPUs: graphics APIs (e.g., DirectX and OpenGL) and computing APIs (e.g., CUDA and OpenCL). First, the graphics APIs provide functions for graphics operations, such as projection, shading, and texture mapping. Second, the computing APIs provide functions for non-graphics applications, such as financial, medical, or weather data analyses, database query optimizations, packet routing, intrusion detection systems, and cryptographic engines.

The most significant differences between the graphics APIs and the computing APIs are sharing and memory manageability. First, the computing APIs allow different users to share the same GPU, whereas the graphics APIs only support a single user. A number of users can share the same GPU using the computing APIs in a time-sharing fashion, as (1) the computing APIs demand no dedicated screens and (2) current GPUs only support sequential execution of different GPU processes. Although some techniques (e.g., VirtualGL) allow remote users to share the same GPU when using the graphics APIs, they warn users of potential security problems (e.g., logging keystrokes and reading back images through an X server). Second, while GPU drivers manage GPU memory with the graphics APIs, programmers can manually manage GPU memory with the computing APIs, including allocations, CPU-GPU data transfers, and deallocations. GPUs have several types of memory (e.g., global, local, and private memories), and programmers can control them using the computing APIs except some graphics-related memories (e.g., framebuffer and z-buffer). In contrast, the graphics APIs provide no functions to manage such memories while providing a set of optimized functions to perform memory-efficient graphics operations.

**Unfortunately,** the sharing and high memory manageability of the computing APIs may incur critical security threats because GPUs do not initialize newly-allocated memory buffers. Although numerous studies consider such an uninitialized memory problem in operating systems, no study deals with the uninitialized GPU memory problem. If similar security threats exist with the computing APIs, the threats have much larger impact as multiple users share the same GPU.

**In this paper,** we first perform an in-depth security analysis on GPUs regarding their architectures and computing APIs to reveal any potential security threats. We identify that the computing APIs have a serious uninitialized memory problem because they (1) do not clear newly allocated memory pages, (2) have memory types that programmers cannot delete, and (3) have in-core memory without security mechanisms.

Second, we develop effective security attacks on GPUs applicable to the most widely used GPUs, NVIDIA and AMD GPUs, by exploiting the revealed security threats. Our attacks can disclose sensitive data kept in GPU memory of a victim program both during its execution and right after its termination.

Third, we demonstrate the high applicability of our attacks by inferring browsing history of the two most widely used web browsers, the Chromium and Firefox web browsers. Both browsers support GPU-accelerated webpage rendering acceleration, which uploads webpage textures to GPU memory to increase rendering speed. Our attacks can extract rearranged webpage textures of both browsers from NVIDIA and AMD GPUs.

-------------

So, yes, we look under yet another stone and discover some security or privacy issues that hadn’t occurred to anyone before. 2019 promises to be another very busy and interesting year!

~30~