Security Now! #634 - 10-24-17

IoT Flash Botnets

This week on Security Now!

This week we discuss some ROCA fallout specifics, an example of PRNG misuse, the Kaspersky Lab controversy, a DNS security initiative for Android, another compromised download occurrence, a browser-based cryptocurrency miner for us to play with... and Google considering blocking them natively, other new protections coming to Chrome, an update on Marcus Hutchins, Microsoft's "TruePlay" being added to the Win10 fall creators update, some interesting "Loopback" from our terrific listeners... and then we take a closer look at the rapidly growing threat of IoT-based "Flash Botnets."

Our Picture of the Week

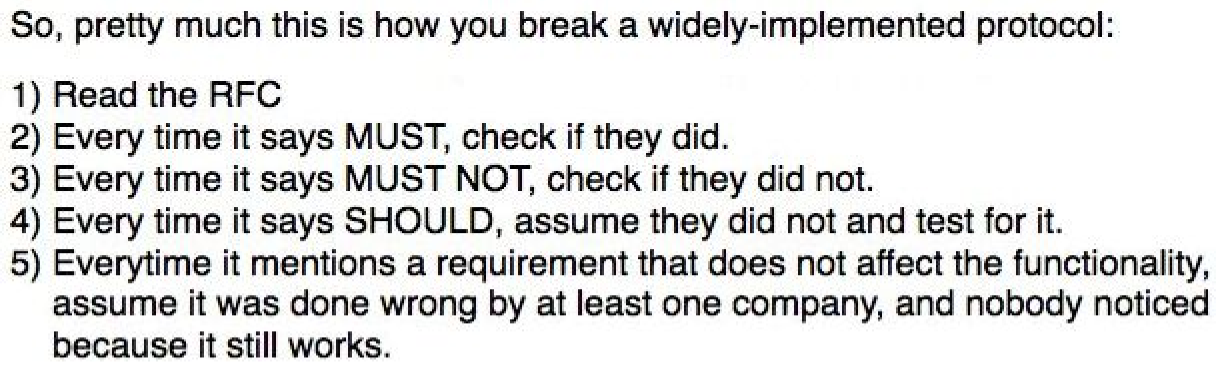

So, pretty much this is how you break a widely-implemented protocol:

1) Read the RFC
2) Every time it says MUST, check if they did.
3) Every time it says MUST NOT, check if they did not.
4) Every time it says SHOULD, assume they did not and test for it.
5) Everytime it mentions a requirement that does not affect the functionality, assume it was done wrong by at least one company, and nobody noticed because it still works.

Security News

ROCA fallout specifics begin to roll in.
https://arstechnica.com/information-technology/2017/10/crippling-crypto-weakness-opens-millions-of-smartcards-to-cloning/

Last week the news had just broken and its impact was still be assessed.

Gemalto has been selling defective Infineon-based smartcards since 2004.

Gemalto won't disclose how many IDPrime.NET cards they have shipped, but it is believed to be as many as hundreds of millions during this time.
The Smartcard was discontinued at the end of last month, September.

However, third-party distributors continue to sell the defective cards online.

In ArsTechnica reporting on this, a Gemalto representative referred to a company advisory, reading: "Our investigation has determined that End-of-sale IDPrime.NET products may be affected."

Cryptography experts added that there is little doubt the line of Gemalto cards is affected. The CEO of Enigma Bridge said he examined eleven of Gemalto's IDPrime.NET cards issued from 2008 through earlier this year. ALL OF THEM used an underlying public key that tested positive for the crippling weakness.

Due to the popularity and ubiquity of these cards, the threat posed is incredibly widespread and will be impossible to completely mitigate. This is the danger imposed by the use of widespread and defective consumer crypto. They will continue to be used and targeted attacks will almost certainly occur... which is exactly what they were intended to prevent.

The government of Estonia has said that the 750,000 (three quarters of a million) electronic IDs it has issued are vulnerable, and researchers have uncovered evidence that the ID cards issued by Slovakia and Spain may be vulnerable, too.

As we knew, several models of TPM from a variety of manufacturers are also known to be affected, as are Javacards.

"Don't Use Hardcoded Keys" -- DUHK -- The DUHK Attack
https://duhkattack.com/
https://duhkattack.com/paper.pdf (Really wonderful paper!)

Two researchers at the University of Pennsylvania and the Johns Hopkins University's prolific Matthew Green uncovered a significant implementation error in at least 12 commercial VPN implementations.

Abstract — The ANSI X9.17/X9.31 random number generator is a pseudorandom number generator design based on a block cipher and updated using the current time. First standardized in 1985, variants of this PRNG design were incorporated into numerous cryptographic standards over the next three decades. It remained on the list of FIPS 140-1 and 140-2 approved random number generation algorithms until January 2016. The design uses a static key with the specified block cipher to produce pseudo-random output. It has been known since at least 1998 that the key must remain secret in order for the random number generator to be secure. However, neither the FIPS 140-2 standardization process in 2001 or NIST’s update of the algorithm in 2005 appear to have specified any process for key generation.

We performed a systematic study of publicly available FIPS 140-2 certifications for hundreds of products that implemented the ANSI X9.31 random number generator, and found twelve whose certification documents use of static hard-coded keys in source code, leaving them vulnerable to an attacker who can learn this key from the source code or binary. In order to demonstrate the practicality of this attack, we develop a full passive decryption attack against FortiGate VPN gateway products using FortiOS version 4. Private key recovery requires a few seconds of computation.We measured the prevalence of this vulnerability on the visible Internet using active scans and find that we are able to recover the random number generator state for 21% of HTTPS hosts serving a default Fortinet product certificate, and 97% of hosts with metadata identifying FortiOSv4. We successfully demonstrate full private key recovery in the wild against a subset of these hosts that accept IPsec connections.

DUHK (Don't Use Hard-coded Keys) is a vulnerability that affects devices using the ANSI X9.31 Random Number Generator (RNG) in conjunction with a hard-coded seed key.

This allows for a "state recovery attack" on a PRNG.

The ANSI X9.31 RNG is an algorithm that has commonly been used to generate the cryptographic keys that secure VPN connections and web browsing sessions.

As we know, cryptography still needs random secrets: "Nonces"

The DUHK allows attackers to recover secret encryption keys from vulnerable implementations and to decrypt and read communications passing over VPN connections or encrypted web sessions.

The researchers uncovered twelve different implementations where traffic can be decrypted by a passive network adversary who is able to simply observe the encrypted handshake traffic. And other key recovery attacks on different protocols may also be possible.

The researchers found a total of twelve FIPS-certified (Federal Information Processing Standard) implementations that document hard-coded X9.31 RNG seed keys in their products and list those in their research.

So... what's the problem? First, here are the requirements. A device is vulnerable to DUHK attacks if:

- It uses the ANSI-standard X9.31 random number generator

and

- The seed key used by the generator is hard-coded into the implementation

and

- The output from the random number generator is directly used to generate keys

and

- At least some of the random numbers before or after those used to make the keys are transmitted unencrypted. This is typically the case for SSL/TLS and IPsec.

<quote>

> A critical design element of the ANSI X9.17/X9.31 PRG is that the cipher key used with the block cipher remains fixed through each iteration. In order to remain secure, the key must never be revealed to external attackers. If the key should become known, an attacker can use the key to decrypt the output and recover all future and past states of the random number generator by brute forcing the timestamp.
>
> Our results show that a non-trivial subset of vendors use static hard-coded keys in source code, leaving them vulnerable to an attacker who can learn this key from the source code or binary.
>
> We are able to perform full state recovery in under a second from random number generator output.

If the PRNG was being used, for example, to generate candidates for primality tests, then its output could never be directly observed and this use would be secure since the isolated primes are never observed alone, only after multiplication, and we know they cannot be readily factored.

But applying this generator in any application where its output can be observed, and where its keying key is also known, allows its internal state to be reverse engineered, and for it to be algorithmically run both forward and backward in time.

Kaspersky Lab to open software to review, says nothing to hide
http://www.reuters.com/article/us-usa-security-kaspersky-russia/kaspersky-lab-to-open-software-to-review-says-nothing-to-hide-idUSKBN1CS0Y1

"We have nothing to hide" - so will open their software we external review and verification.

Kaspersky is in use on 400 million computers worldwide.

In their announcement, Kaspersky did not name the outside reviewers, but said they would have strong software security credentials and be able to conduct technical audits, source code reviews and vulnerability assessments as well as examine Kaspersky's business practices and software development methodology.

Kaspersky said it would open “transparency centers” in Asia, Europe and the United States where customers, governments and others can access results of the outside reviews and discuss any concerns about the security of Kaspersky products.

And they will be expanding their bug bounty program from a maximum award of $5,000 to $100,000.

Not everyone's concerns were mollified by this... but then, nothing Kaspersky could do would make any difference to those who refuse to accept any solution.

Google to add "DNS over TLS" security feature to Android OS
https://www.xda-developers.com/android-dns-over-tls-website-privacy/
https://thehackernews.com/2017/10/android-dns-over-tls.html
https://www.engadget.com/2017/10/23/google-android-dns-tls/
https://developers.google.com/speed/public-dns/docs/dns-over-https

Four days ago, on October 20th, some guys over at XDA Developers (a huge mobile platform developer community) noticed something interesting in the Android Open Source Project (AOSP) software commits.

They wrote: "It appears that “DNS over TLS” support is being added to Android, according to several commits added to the Android Open Source Project (AOSP). The addition in the Android repository shows that a new setting will be added under Developer Options allowing users to turn on or off DNS over TLS. Presumably, if such an option is being added to Developer Options, then that means it is in testing and may arrive in a future version of Android such as version 8.1."

DNS was never designed for security or for privacy.

In an era of extreme spoofing and phishing, DNS's startlingly absent security and privacy is an increasing problem.

DNS is UDP by default for most operations. It can run over TCP, and a few DNS management operations must... but TCP doesn't provide any improvement. It only allows for larger block transfers than UDP.

DNSSEC prevents forgery by having the DNS server provide signatures for its responses. But DNSSEC was never designed to provide privacy.

Last May of 2016, RFC7858 was finalized: "Specification for DNS over Transport Layer Security (TLS)"
https://tools.ietf.org/html/rfc7858

DNS: UDP or TCP over port 53. TLSDNS: port 853.

Connection setup overhead.

Remember: TCP (the underlying transport protocol) has zero bandwidth cost after handshaking.

Since most clients have a "relationship" with only one or two DNS servers, this model works and scales.

And DNS accesses tend to be highly "bursty".

What about authentication and MITM attacks?

Either non-authenticated "opportunistic privacy" or out-of-band Key-Pinned privacy.

Remember though, that today virtually no DNS servers support this very recent specification. But we need to start somewhere. And since this reuses already existing libraries, adding it to the Internet's updated DNS servers should be a relatively simple and straightforward task.

Another instance of a good developer's download servers being compromised:
https://www.eltima.com/blog/2017/10/elmedia-player-and-folx-malware-threat-neutralized.html

We are seeing the emergence of what is known as "supply chain attacks".

In early May a download mirror for the popular HandBrake media converter was compromised to download a Mac malware known as "Proton".

Then the same thing happened with CCleaner.

And now, last Thursday, researchers at ESET discovered that Eltima's "Elmedia Player" for MacOS was also delivering the Proton remote access Trojan (RAT) to unsuspecting and certainly unwanting users. And as before, Eltima's own download server was compromised. Eltima's "Folx" download manager also turned out to be compromised.

Proton first appeared last year and includes many features, such as the ability to execute console commands, access the user’s webcam, log keystrokes, capture screenshots and open SSH/VNC remote connections. It is also able to inject malicious code into the user’s browser to display popups asking victims’ information such as credit card numbers, login credentials, and others. It can also hack into a victim’s iCloud account, even if two-factor authentication is used, and in March of this year it was offered for sale for $50,000 on cybercrime forums.

ESET reported that Eltima was extremely responsive and immediately fixed the problem when notified.

Eltima's own announcement and remediation page provides users with system check information:

SYSTEM CHECK!!!
If you recently downloaded Elmedia Player or Folx, ESET advises you do a system check to confirm if your system was compromised or not.

Instructions: Scan for the absence of the following file or directory on your system:

```
/tmp/Updater.app/
/Library/LaunchAgents/com.Eltima.UpdaterAgent.plist
/Library/.rand/
/Library/.rand/updateragent.app/
```

The presence of any of the files above is an indication that your system may have been infected by the trojanized Elmedia Player or Folx application which means your OSX/Proton is most likely running. If you downloaded Elmedia Player or Folx on the 19th of October 2017, your system is likely affected.

Steps to rid your system of this Malware

A total system OS reinstall is the only guaranteed way to totally rid your system of this Malware. This is a standard procedure for any system compromise with the affection of administrator account.

A sample CryptoCurrency Miner
https://donateyourtab.to/

- Help Puerto Rico (the Hispanic federation)
- Fight Breast Cancer (the breast cancer reearch foundation)
- The Rohingya (considered to be the world's most persecuted minority in Myanmar)
- Civil Rights (the Southern poverty law center)

My creaky old quad-core something-or-other can do about 9 hashes/second at MAX. Many people are reporting about 10x that and comparing one’s various browsers on the same hardware is interesting.

FAQ: How does the mining work?

We use a service called Coin Hive, which is an API that allows browser based cryptocurrency mining. All of the hashes you compute in-browser go into their pool and once we reach a certain dollar amount Coin Hive sends us the money. We then cut a check to each charity based on how much was generated. We make ZERO dollars off this. Coin Hive takes 30% of all proceeds for running the mining pool and providing the mining services. We're exploring ways to make that number lower by using a different or homegrown solution. We also use money raised to cover fees to run this site (hosting, taxes, etc.)

Google Chrome May Add a Permission to Stop In-Browser Cryptocurrency Miners
https://www.bleepingcomputer.com/news/google/google-chrome-may-add-a-permission-to-stop-in-browser-cryptocurrency-miners/

After receiving a number of complaints as "bug reports," an engineer on the Chrome team working on the Chromium project, wrote in one of the recent bug reports:

> "If a site is using more than XX% CPU for more than YY seconds, then we put the page into "battery saver mode" where we aggressively throttle tasks and show a toast [notification popup] allowing the user to opt-out of battery saver mode. When a battery saver mode tab is backgrounded, we stop running tasks entirely.
>
> I think we'll want measurment to figure out what values to use for XX and YY, but we can start with really egregious things like 100% and 60 seconds.
>
> I'm effectively suggesting we add a permission here, but it would have unusual triggering conditions [...]. It only triggers when the page is doing a likely bad thing.

The discussion is ongoing and there's been no consensus reached, but Google is being proactive.

Until then... Chrome users can block in-browser miners via extensions like AntiMiner, No Coin, and minerBlock. And as we mentioned previously, some ad blockers (AdBlock Plus and uBlock Origin) and some A/V products can also block some of these miners.

And what's VERY cool is that with https://donateyourtab.to/ we can test the operation of these various mitigations.

Enable Google's New "Advanced Protection" If You Don't Want to Get Hacked
https://thehackernews.com/2017/10/google-advanced-protection.html

For Chrome on Windows, Google has partnered with ESET to beef up their inadvertently-downloaded malware detection and removal under Chrome's "Cleanup" feature.

<quote>

> Under the hood, we upgraded the technology we use in Chrome Cleanup to detect and remove unwanted software. We worked with IT security company ESET to combine their detection engine with Chrome’s sandbox technology. We can now detect and remove more unwanted software than ever before, meaning more people can benefit from Chrome Cleanup. Note this new sandboxed engine is not a general-purpose antivirus—it only removes software that doesn’t comply with our unwanted software policy.
>
> We’ve begun to roll this out to Chrome for Windows users now. Over the next few days, it will help tens of millions of Chrome users get back to a cleaner, safer web.

Marcus Hutchins (aka MalwareTech) may no longer be under curfew or GPS monitoring
Marcus Hutchins, awaiting trial, can now live and work unencumbered in LA.
https://arstechnica.com/tech-policy/2017/10/judge-malwaretech-is-no-longer-under-curfew-gps-monitoring/

Last Thursday, one of Marcus' attorneys successfully argued that these restrictions were needlessly burdensome. He had never missed a court appearance, his GPS device failed during a trip East and he didn't attempt to flee, and he wants to swim and surf which are impeded by the need to wear the GPC tracker. Also, Marcus' Internet access restrictions have been lessened.

Marcus' Tweet: "Californians claiming it's cold meanwhile I'm wishing there was a way to wear less clothes than I'm wearing without being arrested. — MalwareTech (@MalwareTechBlog) October 20, 2017

However... by the end of the next day the DoJ had filed a "Motion to Revoke" the judge's decision.

Michael Chmelar, Assistant United States Attorney: "While it is true pre-trial services possesses defendant's passport, it is unrealistic to think that the defendant could not leave the U.S. without travel documents."

(Marcus later tweeted that he was still in limbo.)

Windows 10 Fall Creator's Update adds "TruePlay" - anti-Game Cheating technology.
(Because... Microsoft had some spare time on their hands??)
https://arstechnica.com/gaming/2017/10/microsoft-rolls-out-system-level-anti-cheating-tech-for-windows-devs/
https://msdn.microsoft.com/en-us/library/windows/desktop/mt808781(v=vs.85).aspx

ArsTechnica: "Lexicon, Force Hax, and Menyoo were all subscription-based paid hacking tools that let GTA Online players spawn infinite piles of cash, teleport other players to arbitrary locations, become invulnerable, or walk through walls while playing with other people. Over the weekend, though, the websites for all three programs were replaced with a simple message:

" After discussions with Take-Two Interactive, effective immediately we are ceasing all maintenance, development and distribution of [our] cheat menu services. We will be donating our proceeds to a charity designated by Take-Two. We apologize for any and all problems [our program] has caused to the Grand Theft Auto Online community. "

Microsoft's forthcoming "TruePlay" technology allows game developers to easily mark their titles are "high protection" which causes Windows 10 FCU to enhance the process' isolation to significantly increase the difficulty of interfering with the local and networked game play.

This form of process isolation has been possible and available for developers who wished to implement it themselves. But Microsoft's incorporation of this into Windows will make its application more universal... and it allows Microsoft and Windows 10 to evolve this protection over time.

SpinRite

Rich Williams
Location: Daphne, AL
Subject: One more for SpinRite
Date: 23 Oct 2017 15:27:04
:
Hi Steve,

I know you receive a lot of feedback and have discussed SpinRite and its ability to repair many storage devices, but I ran into something last week you (and listeners) may want to know. My Samsung S7 (which runs Android) has been having a problem recently where I would regularly be notified that my SD card was Encrypted. This is normal at boot and everything seemed to run fine, so I didn't worry about it at first, but after some research to stop the annoyance, I found that it was being caused by an unmount/remount operation happening when Android failed to read the card. I store my podcasts, pictures and videos on it and it holds 256GB so I can't afford to start losing data gradually. Needless to say, I pulled out my copy of SpinRite and after it completed running on level 2 with no errors... It's been 3 days now and no more alerts.

Thank you for a great product, and for the work you do on Security now.

Best Regards,
Rich

Loopback:

Domenico Lamberti (@Mobile_Dom)
I gotta say,after hearing @leolaporte and @SGgrc 's take on mining in browsers, I'm totally for sites using JavaScript miners for revenue.

(((@dcliterate))) (@DCLiterate)
I am sure you know about Google Advanced Protection by now @SGgrc
My Q: does mobile use require re auth w/ physical key each time?

Simon Zerafa (@SimonZerafa)
@SGgrc FYI: Mail Store Home 10.2 is now available:
https://www.mailstore.com/en/products/mailstore-home/
https://www.mailstore.com/en/products/mailstore-home/changelog/

Ronald Collins (@reclusecannon)
@SGgrc I'm a big e-ink fan, thinking of jumping from original Kindle Paperwhite to new model Oasis...does it get your seal of approval?

Anthony Carter @realfoodman
@SGgrc Question for podcast: with so much IPv6 space, is it theoretically possible to assign a unique IP to every process on a computer?

Kyle Hardin (@kylehardin_AK)
@SGgrc RE: mobile sign in spoofing, seems the issue is that services often ask you to prove who you are before they prove who they are.

James Wilson (@sage123)
@SGgrc TU for the great podcast. If a site loses your credentials do you need to reset your oauth 2fa?

The Clear and Present Danger of Flash IoT Botnets

"IoT Reaper" aka "IoTroop"

A Gigantic IoT Botnet Has Grown in the Shadows in the Past Month

- https://www.bleepingcomputer.com/news/security/a-gigantic-iot-botnet-has-grown-in-the-shadows-in-the-past-month/
- https://thehackernews.com/2017/10/iot-botnet-malware-attack.html
- https://threatpost.com/iotroop-botnet-could-dwarf-mirai-in-size-and-devastation-says-researcher/128560/

The Hacker News: “New Rapidly-Growing IoT Botnet Threatens to Take Down the Internet”
Bleeping Computer: “A Gigantic IoT Botnet Has Grown in the Shadows in the Past Month”
ThreatPost: “‘IOTroop’ Botnet Could Dwarf Mirai in Size and Devastation, Says Researcher”

It’s been exactly a year since “Mirai” (Japanese for “the future”), the largest incidence of IoT-based malware seen before now, wreaked havoc by bringing down and holding down the DynDNS services long enough for the Internet’s DNS caches to expire and render many Dyn-serviced sites inaccessible. Once upon a time we could have substituted a site’s IP address, but not in today’s world of HTTPS which requires a domain name match with the remote site’s certificate.

Now, on the occasion of this first anniversary, the Internet is facing a new and much greater rapidly-emerging threat. A newly detected IoT botnet is sweeping the Internet, and sweeping up many defective IoT devices… putting millions of IoT devices under its control.

Whereas Mirai scanned for open Telnet ports and attempted to log into devices using a preset list of default or weak administrative credentials, this newly discovered botnet -- named either “IoT Reaper” or “IoTroop” by its two separate discoverers -- is spreading rapidly across the Internet by scanning for exploitable and known but unpatched vulnerabilities in popular routers, cameras and network video recorders by major brand manufacturers:

- Dlink (routers)
- Netgear (routers)
- Linksys (routers)
- Goahead (cameras)
- JAWS (cameras)
- AVTECH (cameras)
- Vacron (NVR)

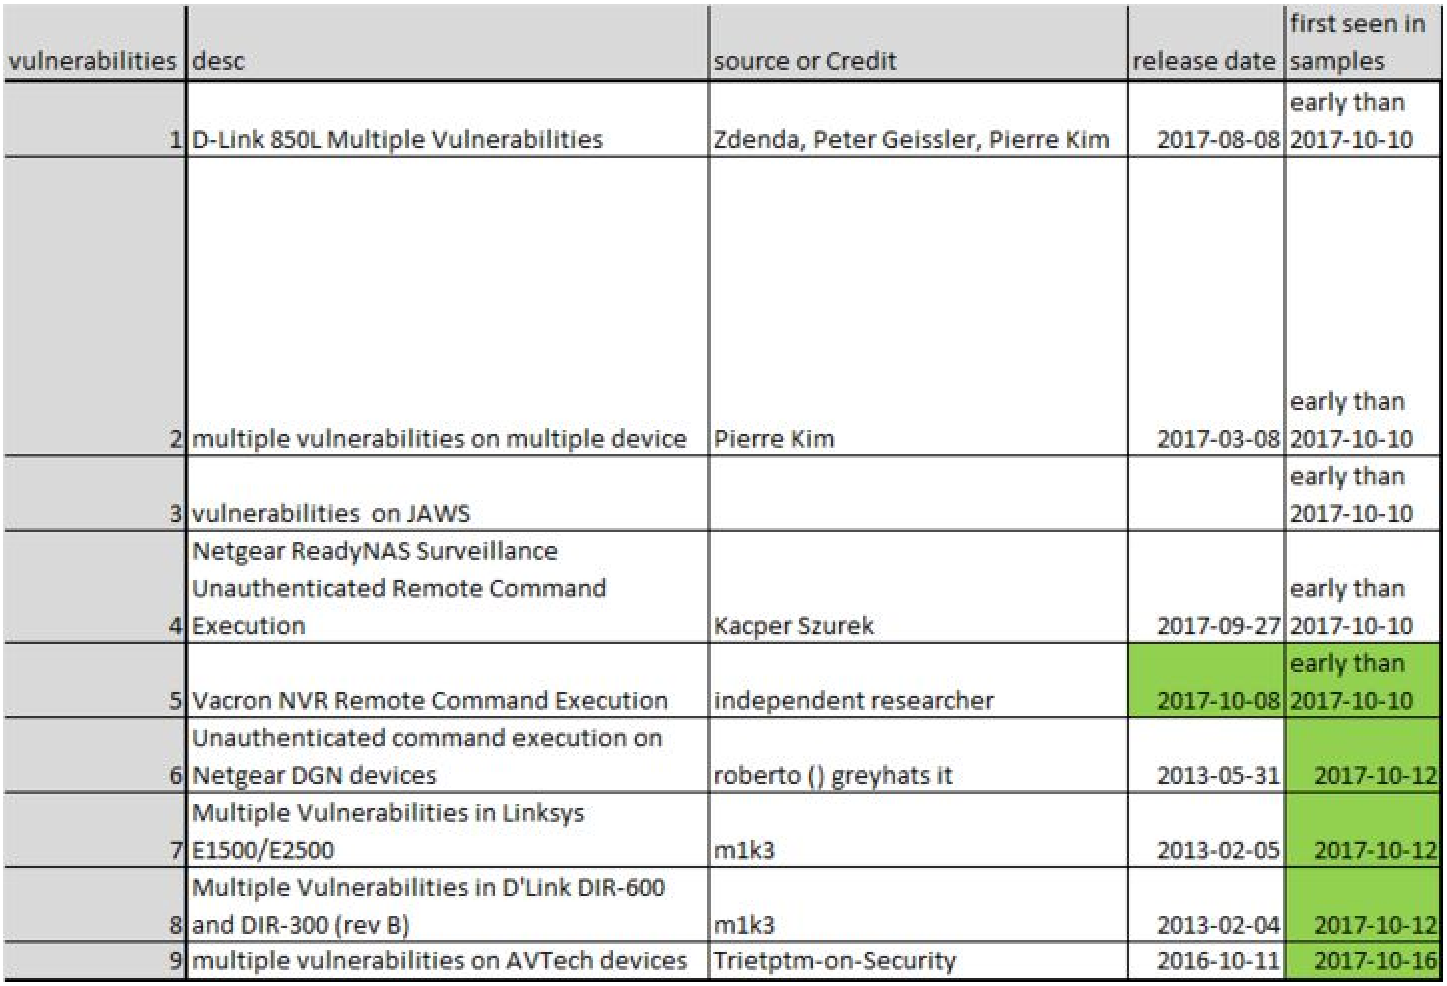

vulnerabilities	desc	source or Credit	release date	first seen in samples
1	D-Link 850L Multiple Vulnerabilities	Zdenda, Peter Geissler, Pierre Kim	2017-08-08	early than 2017-10-10
2	multiple vulnerabilities on multiple device	Pierre Kim	2017-03-08	early than 2017-10-10
3	vulnerabilities on JAWS			early than 2017-10-10
4	Netgear ReadyNAS Surveillance Unauthenticated Remote Command Execution	Kacper Szurek	2017-09-27	early than 2017-10-10
5	Vacron NVR Remote Command Execution	independent researcher	2017-10-08	early than 2017-10-10
6	Unauthenticated command execution on Netgear DGN devices	roberto () greyhats it	2013-05-31	2017-10-12
7	Multiple Vulnerabilities in Linksys E1500/E2500	m1k3	2013-02-05	2017-10-12
8	Multiple Vulnerabilities in D'Link DIR-600 and DIR-300 (rev B)	m1k3	2013-02-04	2017-10-12
9	multiple vulnerabilities on AVTech devices	Trietptm-on-Security	2016-10-11	2017-10-16

The Vacron NVR exploit was added to IoT Reaper within 48 hours of its public disclosure. So someone is watching and actively evolving this botnet rapidly.

Check Point also spotted the botnet attacking MicroTik adn TP-Link routers, Synology NAS devices, and Linux servers.

The malware also also integrates a full LUA execution environment, allowing the author to write very complex and efficient attack scripts.

Two teams of researchers who have been monitoring and reverse-engineering the botnet’s operation believe that this latest “IoT Reaper” malware malware has already infected nearly ***two million devices*** and is growing continuously at the extraordinary rate of ~10,000 new devices per day!

To put the attack power of this into perspective, Mirai was able to take DynDNS down with 5% of that count -- approximately 100,000 devices.

And the researchers who have now examined this botnet's code found that it was aware of more than 100 DNS open resolvers which would enable it to launch powerful DNS amplification attacks.

And, perhaps saddest of all, this is not a new problem. In early April of 2014 -- so just over three and a half years ago -- The Hacker News posted a story with the headline: "Millions of Vulnerable Routers aiding Massive DNS Amplification DDoS Attacks"
https://thehackernews.com/2014/04/millions-of-vulnerable-routers-aiding_3.html

And among those two million infected devices are at least 100,000 enterprises whose internal security, while not the apparent present target of the botnet, is entirely at risk while their devices are compromised.

At this time no one knows who created this and why, but the DDoS threat landscape is skyrocketing and could reach tens of terabits-per-second in size.

Researchers at CheckPoint wrote: "Our research suggests we are now experiencing the calm before an even more powerful storm. The next cyber hurricane is about to come."

So… WHERE does this leave us?

~30~