Security Now! #635 - 10-31-17

Reaper Redux

This week on Security Now!

This week we examine the source of WannaCry, a new privacy feature for Firefox, Google's planned removal of HPKP, the idea of visual objects as a second factor, an iOS camera privacy concern, the CAPTCHA wars, a horrifying glimpse into a non-Net Neutrality world, the CoinHive DNS hijack, the new Bad Rabbit crypto malware, a Win10 anti-crypto malware security tip, spying vacuum cleaners, a new Amazon service, some loopback Q&A with our listeners and another look at the Reaper botnet.

Our Picture of the Week

Security News

North Korea appears to have been behind WannaCry
https://www.nytimes.com/2017/10/27/world/europe/uk-ransomware-hack-north-korea.html

As we know, reliable attribution for cyberattacks is made difficult by the ability for attacks to be looped through and bounced off of remotely located innocent machines. But the UK and other nations have arrived at the consensus that the now-famous WannaCry cyberattacks of last May 12th, which crippled the UK's national health services and spread to more than 150 countries by leveraging a flaw in Windows' SBMv1... were designed and initiated by North Korea's 6,000 strong hacker regime.

Firefox v58 will be the first widely used browser to adopt anti-canvas fingerprinting technology
https://www.bleepingcomputer.com/news/software/firefox-implements-another-privacy-preserving-feature-taken-from-the-tor-browser/
https://nakedsecurity.sophos.com/2017/10/30/firefox-takes-a-bite-out-of-the-canvas-super-cookie/

We've discussed "browser fingerprinting" in detail.
https://panopticlick.eff.org/

There are many static beacons which JavaScript is able to query:

- Browser plug-ins
- System fonts (Blocked in FFv52, also inherited from the TOR browser.)

Drawing text on an off-screen "canvas" surface then hashing the resulting image is surprisingly effective. Different browsers, OS platforms, font versions, and GPU accelerators result in images which are exactly reproducible on a single machine and browser, while differing from system to system:

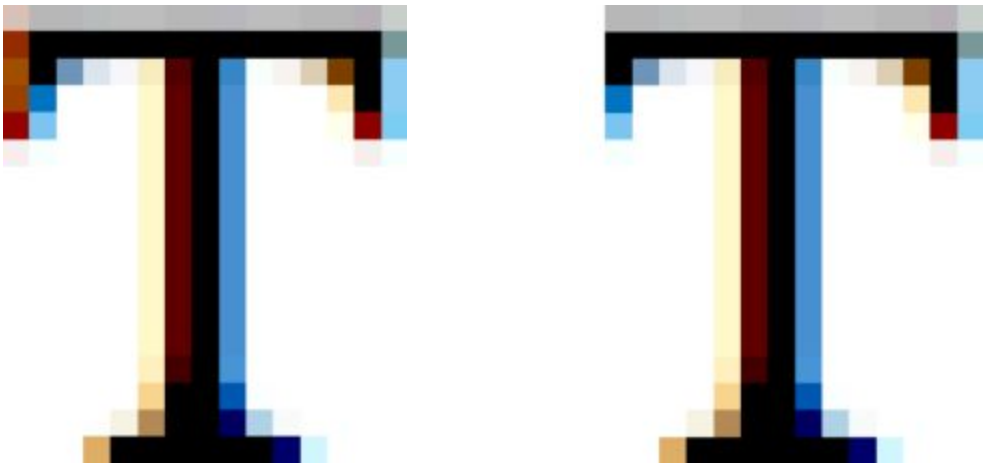

Four years ago, the privacy focused TOR browser added the feature to prompt the user for permission when a web page attempted to extract the image data from the browser canvas.

Today, the mainstream Firefox browser will be inheriting that feature.

Firefox 58 is scheduled for release on January 16, 2018.

Google to Remove Public Key Pinning (PKP) Support in Chrome
https://www.bleepingcomputer.com/news/security/google-to-remove-public-key-pinning-pkp-support-in-chrome/
https://groups.google.com/a/chromium.org/forum/#!msg/blink-dev/he9tr7p3rZ8/eNMwKPmUBAAJ

PKP - Public Key Pinning - RFC7469 (HTTP PKP - HPKP)

Not all ideas are good.

Key pinning is a trust-on-first-use (TOFU) mechanism.

```
Public-Key-Pins: max-age=2592000;
        pin-sha256="E9CZ9INDbd+2eRQozYqqbQ2yXLVKB9+xcprMF+44U1g=";
        pin-sha256="LPJNul+wow4m6DsqxbninhsWHlwfp0JecwQzYpOLmCQ="
```

... but it all turned out to be brittle, fragile, error-prone, and more trouble than it was worth.

It will be gone at Chrome 67 (May 2018)

Objects as a second factor?
https://www.theverge.com/2017/10/26/16553900/2fa-two-factor-authentication-pixie-mobile-devices

Smartphone takes a photo of a "trinket"

A research paper, not meant for implementation.

Replay attack - same static thing every time.

A time-based TOTP is powerful because it cannot be replayed. The danger of TOTP is that it uses a shared secret. This was the huge problem caused by the security breach at RSA back in 2011 which endangered and compromised its SecureID product line.

A public-key challenge/response system (such as SQRL) is even more powerful, since it doesn't reply upon the authenticating server keeping a shared secret.

Too much of a good thing? Granting iOS apps camera permission
https://krausefx.com/blog/ios-privacy-watchuser-access-both-iphone-cameras-any-time-your-app-is-running

Felix Krause: He's the researcher who recently demonstrated how easily malicious app could spoof iOS credential requests.

Once an app has been granted access to your iOS device's camera, it can:

- access both the front and the back camera
- record you at any time the app is in the foreground
- take pictures and videos without telling you
- upload the pictures/videos it takes immediately
- run real-time face recognition to detect facial features or expressions

... all without indicating that your phone is recording you and your surrounding, no LEDs, no light or any other kind of indication.

Felix asks: “Have you ever used a social media app while using the bathroom?”

The application MUST be in the foreground... but it doesn't need to provide any indication that it's actively taking pictures or streaming video from your device's cameras.

Just something to be aware of and to keep in mind.

Time to do a little privacy audit: "Privacy" under "Settings" app.

The unintended consequences of Bots and BotNets.
unCAPTCHA Breaks 450 ReCAPTCHAs in Under 6 Seconds
https://www.bleepingcomputer.com/news/technology/uncaptcha-breaks-450-recaptchas-in-under-6-seconds/

The cat and mouse back-and-forth game of CAPTCHAs has now created the "unCAPTCHA."

CAPTCHA: Completely Automated Public Turing Test To Tell Computers and Humans Apart.

unCAPTCHA is the name of an experimental automated system designed by a team of University of Maryland computer scientists that can break Google's reCAPTCHA challenges with an accuracy of 85%.

The system doesn't target reCAPTCHA's image-based challenges, but the audio version that Google added so people with disabilities can solve its puzzles.

unCAPTCHA downloads the audio puzzle then feeds it to six speech recognition systems: Bing Speech Recognition, IBM, Google Cloud, Google Speech Recognition, Sphinx, and Wit-AI. It aggregates the results and feeds most likely answer back to Google's servers.

In tests performed by the researchers, their experimental system was able to break 450 reCAPTCHA challenges, with 85.15% accuracy, in 5.42 seconds... less time than a person would require to listen to one reCAPTCHA audio challenge.

And on the VISUAL CAPTCHA side...

Last week, other researchers announced they created an AI vision bot that can also break various CAPTCHA systems with high accuracy. This new system solved Google reCAPTCHAs with 66.6% accuracy, BotDetect with 64.4%, Yahoo with 57.4%, and PayPal image challenges with 57.1%.

In Portugal, with no net neutrality, internet providers are starting to split the net into packages.

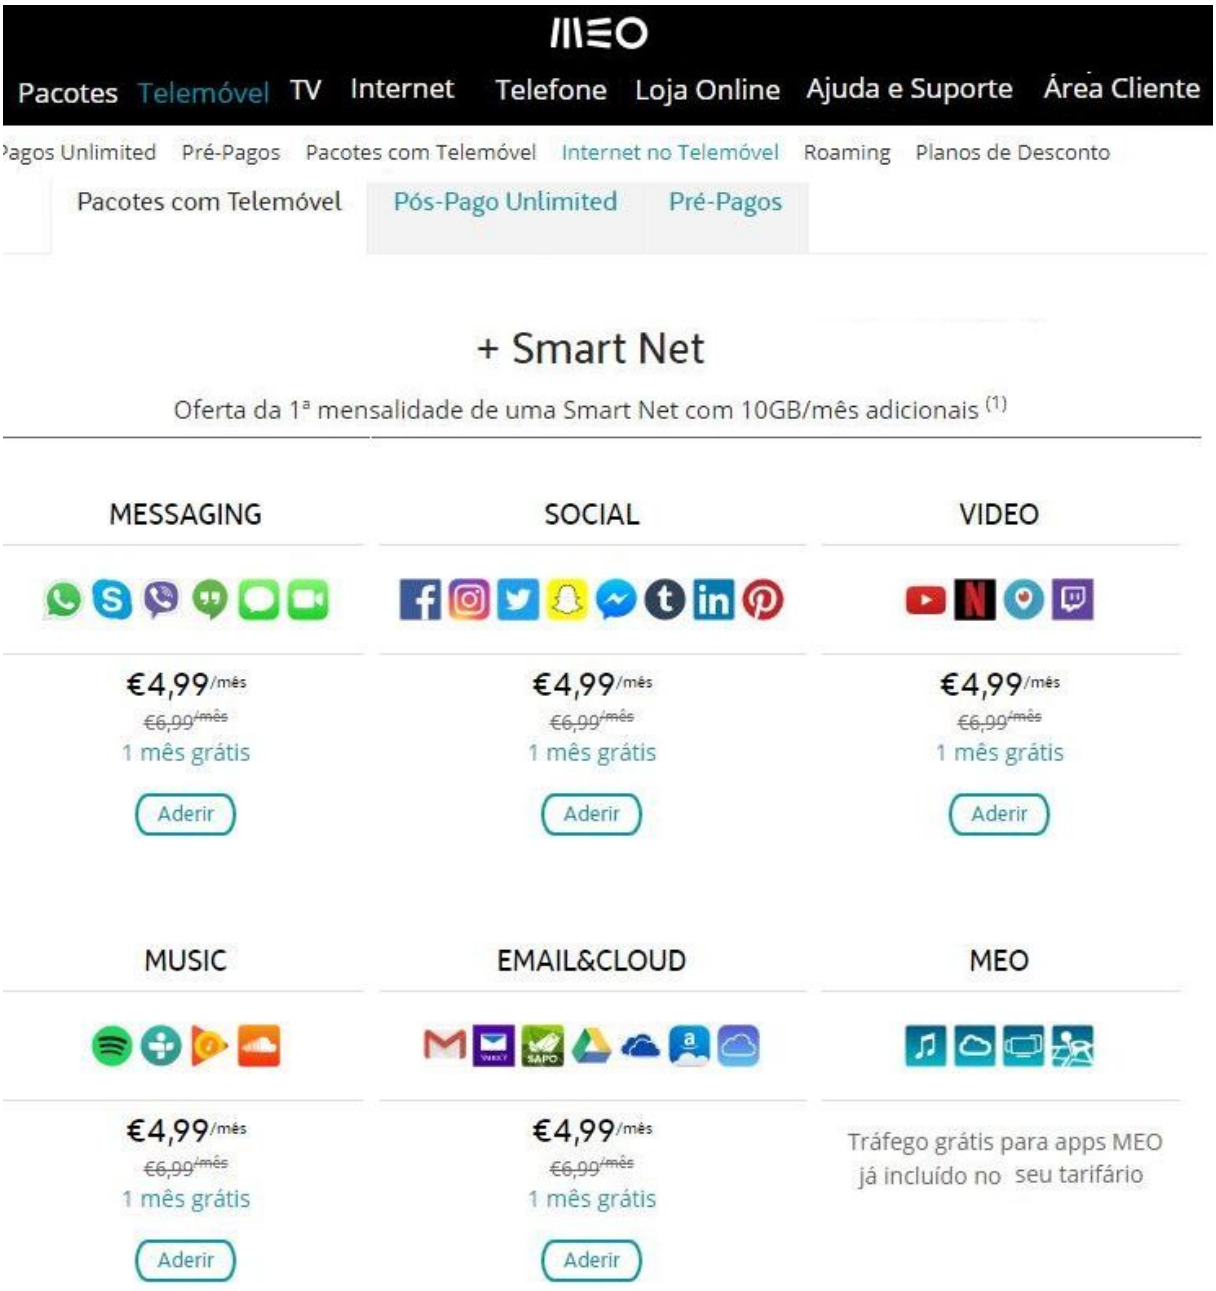

In-Browser Mining Script Provider CoinHive Suffers Major DNS Hijack
https://themerkle.com/in-browser-mining-script-provider-coinhive-suffers-major-dns-hijack/
https://thehackernews.com/2017/10/coinhive-cryptocurrency-miner.html
http://www.securitybreach.online/2017/10/27/coinhive-dns-server-used-mine-monero-crypto-currency-unknown-hacker/

Coinhive is the primary browser-based Monero cryptocurrency hosting site.

Three years ago, back in 2014, a major data breach at Kickstarter exposed a large number of passwords.

When Coinhive setup their hosting with Cloudflare they reused the same password which had been leaked by the Kickstarter breach and neglected to setup any of the other available second-factor authentication.

This allowed attackers to get into their Cloudflare account and commandeer their DNS to point the "coinhive.com" domain to their own servers... which effectively commandeered the entire coinhive crypto currency mining operation in one fell swoop.

```
<script src="https://coinhive.com/lib/coinhive.min.js"></script>
<script>
    var miner = new CoinHive.User('<site-key>', 'john-doe');
    miner.start();
</script>
```

Whoopsie!

Bad Rabbit: New Ransomware Attack Rapidly Spreading Across Eastern Europe
https://thehackernews.com/2017/10/bad-rabbit-ransomware-attack.html

Like WannyCry and several other recent malwares, Bad Rabbit is leveraging the NSA's "EternalRomance" vulnerability released by the Shadow Brokers and subsequently patched by Microsoft last March.

But as we well know, the availability of a patch and the installation of the patch are very different things.

"EternalRomance" is a remote code execution exploit that takes advantage of a flaw (CVE-2017-0145) in Microsoft's Windows Server Message Block (SMB), a protocol for transferring data between connected Windows computers, to bypass security over file-sharing connections, thereby enabling remote code execution on Windows clients and servers.

Bad Rabbit was being distributed by drive-by download attacks via compromised Russian media sites, using fake Adobe Flash players installer to lure victims' into install malware unwittingly and demanding 0.05 bitcoin (~ $285) from victims to unlock their systems.

Once an instance of Bad Rabbit get onto a single machine is scans the internal network for open SMB shares, tries a hardcoded list of commonly used credentials to drop malware, and also uses Mimikatz post-exploitation tool to extract credentials from the affected systems.

It's also able to exploit the Windows Management Instrumentation (WMI) Command-line (WMIC) scripting interface to execute code on other Windows systems on the network remotely.

And Cisco's Talos group weighed in to note that Bad Rabbit also carries a code that uses EternalRomance, which allows remote hackers to propagate from an infected computer to other targets more efficiently.

Our standard advice is:

- To NEVER download anything offered to you by any website. Even trusted sites may have been (and often are) transiently compromised. So it doesn't matter who is making the offer for something your system needs (just like we learned that eMail from our mom's wasn't necessarily safe!)

- Place all networks behind NAT-routing firewalls with UPnP disabled where possible.

- Establish and maintain network isolation wherever possible.

- And... keep up to date with all system patching.

Windows 10 tip: Turn on the new anti-ransomware features in the Fall Creators Update
http://www.zdnet.com/article/windows-10-tip-turn-on-the-new-anti-ransomware-features-in-the-fall-creators-update/

Ed Bott, wroting for ZDNet reminds us that the Win10 Fall Creators Update adds a powerful and apparently effective anti-cryptomalware feature known as "Controlled Folder Access"

Windows Defender must be enabled -- It is NOT enabled by default.

Under Windows Defender Security Center. Click Virus & threat protection > Virus & threat protection settings and flip the switch under Controlled folder access to On.

The feature's default settings may be adjusted using the two links located on the UI:

Click Protected folders to display a list of the folders whose contents are currently being protected from tampering by a malicious or suspicious app. The default list includes data folders from the current user profile and from the Public profile.

Click Allow an app through Controlled folder access to manage a list of whitelisted apps. Most well-known apps are already whitelisted, but you can add a program to the list if you see a Controlled Folder Access error message from that app and know that it's safe and trustworthy.

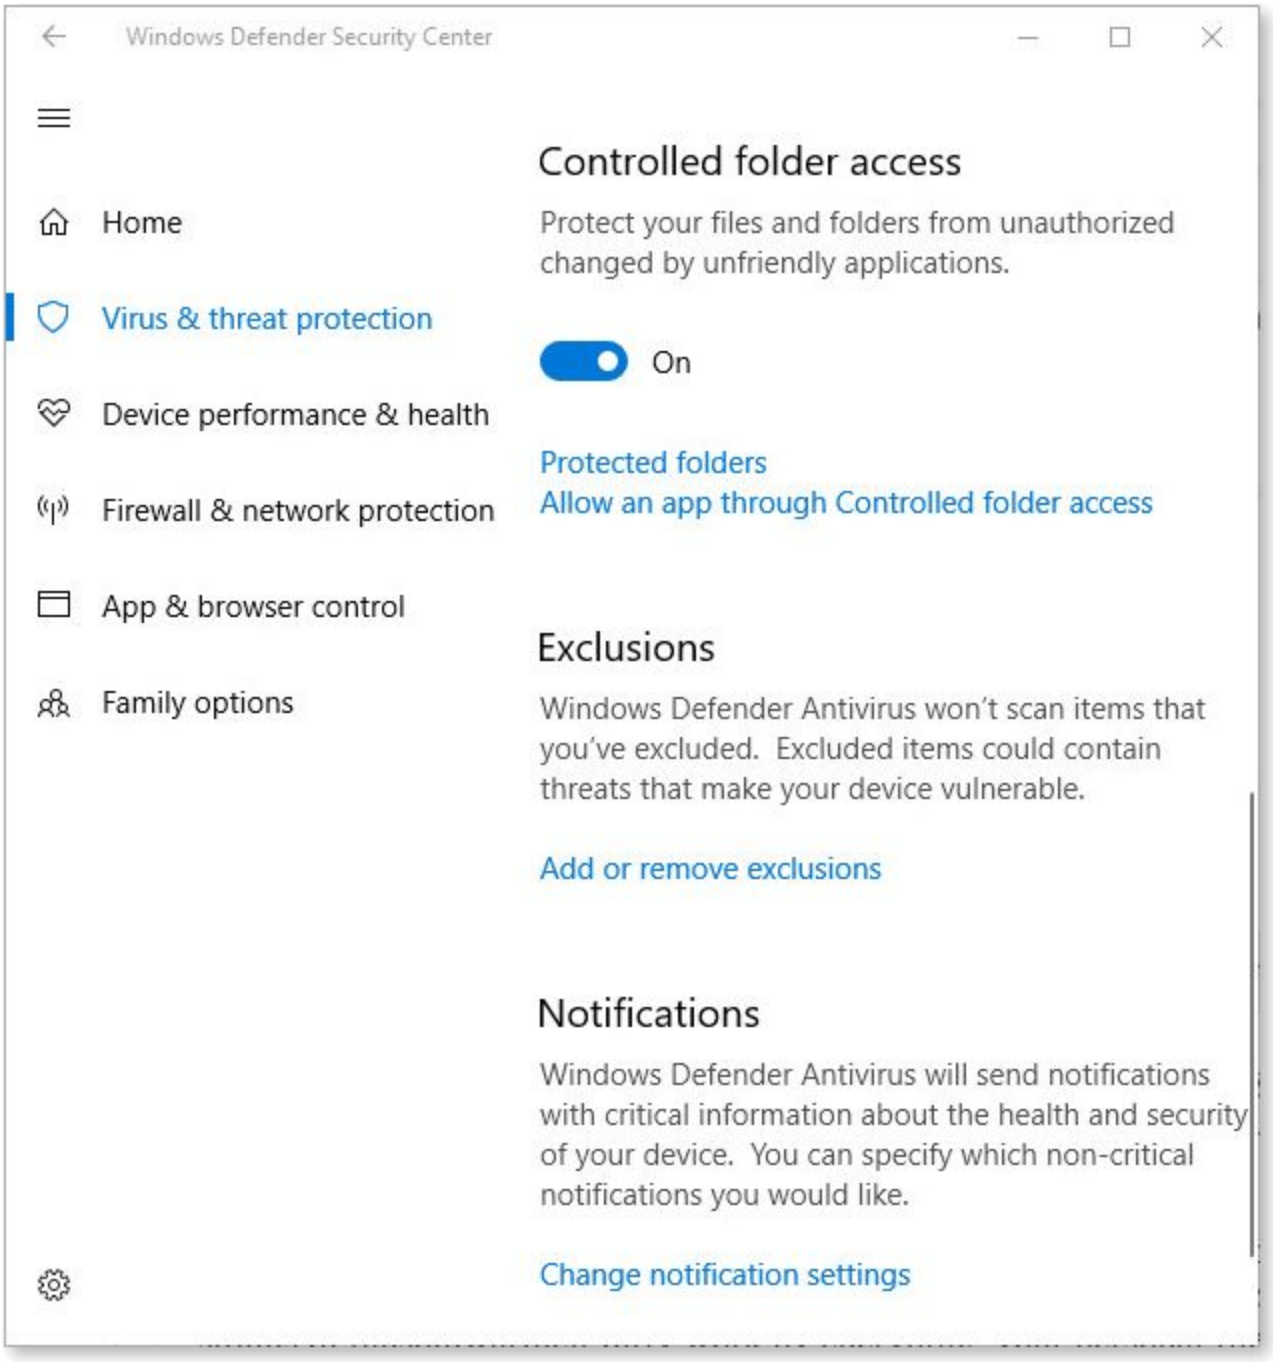

Not only did CheckPoint spot the IoTroop botnet...
but they also found a bad flow in LG's Hom-Bot robot vacuum cleaner.
https://www.digitaltrends.com/home/lg-hom-bot-vacuum-hacked/
https://thehackernews.com/2017/10/smart-iot-device-hacking.html

Check Point researchers discovered a security vulnerability in LG SmartThinQ smart home devices that allowed them to hijack internet-connected devices like refrigerators, ovens, dishwashers, air conditioners, dryers, and washing machines manufactured by LG.

Dubbed HomeHack, this was a vulnerability in LG's smart home infrastructure exposing it to critical user account takeover. If attackers had exploited this vulnerability they would have been

able to log into LG users' SmartThinQ® home appliances accounts and take remote control of the devices connected to the account.

The HomeHack vulnerability gave attackers the potential to spy on users' home activities via the Hom-Bot robot vacuum cleaner video camera (which numbered more than 400,000 by the first half of 2016). The Hom-Bot sends live video to the associated LG SmartThinQ app as part of its HomeGuard Security feature. Depending on the LG appliances in the owner's home, attackers could also switch dishwashers or washing machines on or off.

Checkpoint notified LG about this vulnerability on July 31 2017, and LG responded responsibly to stop possible exploitation of the issues in its SmartThinQ app and devices, releasing a new version patching this vulnerability at the end of September.

It should be becoming quite clear that IoT devices MUST be provided by responsible companies and MUST have autonomous self-patching facility built-in. End users cannot be required or relied upon to patch their vacuum cleaners.

"Amazon Key" service with the new Amazon Cloud Cam
https://www.theverge.com/2017/10/25/16538834/amazon-key-in-home-delivery-unlock-door-prime-cloud-cam-smart-lock

The new service is called Amazon Key, and it relies on a Amazon's new Cloud Cam and a compatible smart lock.

The camera is the hub, connected to the internet via residential Wi-Fi.

The camera talks to locks manufactured by Yale and Kwikset over the Zigbee wireless protocol.

When a courier arrives with a package for in-home delivery, they scan the barcode, sending a request to Amazon's cloud. If everything checks out, the cloud grants permission by sending a message back to the camera, which starts recording. The courier then gets a prompt on their app, swipes the screen, and the homes' front door unlocks.

The package is dropped off and the door is relocked with another swipe.

The customer receives a notification that their delivery has arrived, along with a short video showing the drop-off to confirm everything was done properly.

Who is Satoshi Nakamoto?
https://www.cnbc.com/2017/10/27/bitcoins-origin-story-remains-shrouded-in-mystery-heres-why-it-matters.html

Bitcoin's creator may be worth $6 billion — but people still don't know who it is

Despite some highly publicized false alarms, we still don't know who originated the Bitcoin cryptocurrency. It remains one of the biggest mysteries in the technology world.

But we do have some idea of the inventor's net worth, since "Satoshi" is believed to be holding about five percent of all issued bitcoins having an estimated value of ~$6 Billion.

Tunnelbear
https://www.tunnelbear.com/blog/how-we-avoid-collecting-your-ip-address-on-our-website/

"How We Avoid Collecting Your IP Address On Our Website"

> <quote> It might sound odd, but we pride ourselves on knowing very few things about our customers. In fact, the less we know about you the happier we are. As we grow and learn more about the information we need to optimize our service, we find new ways to do things like anonymize site statistics in a privacy compliant way.
>
> You see, we want to make sure that even before you decide you need a Bear in your life, we're helping you maintain your privacy. Here are a few of the things we do on our website to limit the collection of your IP address.

They then go on to describe how they deliberately use a feature offered by Google Analytics to obscure the last byte of every visitor's IP address, recording it as .0 regardless of its actual IP. In practice, this preserves geographic data but hides an individual ISP customer among 255 others.

Errata

Khyron @crlowell
@SGgrc - Really enjoying #SecurityNow!
Last week you mentioned bitcoin mining via JS, iirc most JS miners are using #Monero, not BTC.

Miscellany

Stranger Things Season #2

SpinRite

From: "Rusty Burke"
Subject: Security Now Series on "Networking"

Before I get to my question (which is not actually sales related), I just want to express my thanks to SpinRite. It recovered a Sony Vaio laptop which was, with increasing frequency, inexplicably dying on me.

So I ran SpinRite at level 4 and after re-booting, I haven't had a lick of trouble. Runs like new.

My question is whether Steve ever did a series on "networking" which he suggested might be coming while doing the "how computers work" series on Security Now. I've combed through several sources of an answer to this question and can't seem to find one.

Thanks for a great product and great information, Rusty Burke

Closing The Loop

Steven @Ozzyla
@SGgrc Just heard on your podcast about running SpinRite on a SD Card, how the hell you do that? Tutorial anywhere?

Robert D Wilson @robertdwilson
@SGgrc Is 'Dark Matter' on @netflix any good?

Jan Rademan @radjanoonan
@SGgrc Is there a test to see if a router is a Reaper zombie? Could Shieldsup be used to test for specific port?
Dr Brian of London @brianoflondon
@SGgrc is there a way to see if my own ISP provided D-link router is part of the new mega bot-net?

warrenkc @warrenkc
@GibsonResearch In your opinion, should a user disable windows defender on Windows 10? If said user is careful on the internet...??Thank You!

Tim Chase @gumnos
@SGgrc SN632 on DNSSEC said cert info could be kept in DNS.
This would only be DV certs, not EV, right?
Grant Taylor @DrScriptt
Replying to @gumnos @SGgrc
Why couldn't an EV cert be stored in DNS via TLSA (DANE)?
Or are you implying DANE *ONLY* /implies/ DV?

Reaper Redux

https://research.checkpoint.com/iotroop-botnet-full-investigation/

https://arstechnica.com/information-technology/2017/10/assessing-the-threat-the-reaper-botnet-poses-to-the-internet-what-we-know-now/

https://f5.com/labs/articles/threat-intelligence/cyber-security/reaper-the-professional-bot-herders-thingbot

Reaper: The Professional Bot Herder's "Thingbot"

Original posting by **F5 Labs**, last Thursday the 26th:

This isn't your mama's botnet. This is a proper botnet. If you were the world's best IoT botnet builder and you wanted to show the world how well-crafted an IoT botnet could be, Reaper is what you'd build. It hasn't been seen attacking anyone yet, and that is part of its charm. But, what is it doing? We've got some ideas.

Oct 31, 2017 Update

The intentions of Reaper are as unclear today as they were a week ago. We hold to our position that the interesting aspect of Reaper is not its current size, but its engineering, and therefore its potential.

From a pure research perspective, we're interested in how Reaper is spreading. Instead of targeting weak auth like a common thingbot, Reaper weaponizes nine (and counting) different IoT vulnerabilities.

We think the current media focus on "the numbers" instead of the method is a tad myopic. See the next "update" section below for our clarification.

Original posting:

Size and Position

Krebs puts the current size of Reaper at over one million IoT devices. We have data that suggests it could include over 3.5 million devices and could be capable of growing by nearly 85,000 devices per day. The reason Reaper has gotten so big and, honestly, the reason we're so impressed with its construction is that, unlike its predecessors, Mirai and Persirai, Reaper uses multiple attack vectors. Mirai used default passwords. Persirai used the blank username + password combo, which frankly is such a doofus security error on the part of the manufacturer that we feel it barely deserves to have a CVE.
Reaper is almost showing off by not even trying the password cracking, and instead just exploiting different vulnerabilities (RCEs, web shells, etc.) in nine different IoT vendor devices.

Oct 31, 2017 Update (continued)

Reports on the "size" of Reaper vary. We've scanned 750,000 unique devices that match the nine vulnerabilities currently exploited by Reaper. We regularly scan 85,000 new, "Reaper-compatible" devices per day. We don't know which of them are actually infected, but there's no reason that Reaper itself couldn't infect them, unless its authors didn't want it to.

The nine vulnerabilities currently used by Reaper are fairly rudimentary, as vulnerabilities go. If the thingbot authors were to include a few dozen existing vulnerabilities that fit Reaper's device-targeting profile, we think they could grow the thingbot by an additional 2.75 million nodes. If they wanted to. Adding that 2.75 million to the 750,000 that are currently "Reaper-compatible" gives the number 3.5 million.

Note: We will not be disclosing the additional CVEs as that would simply expedite the authors' exploits.

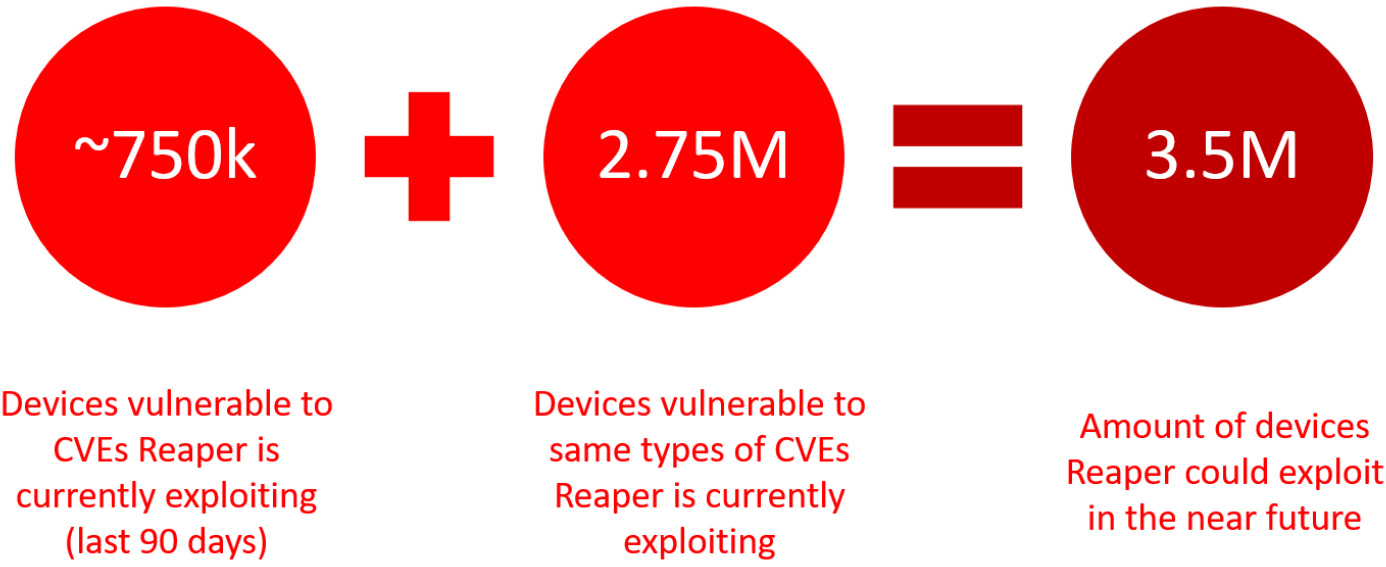

The actual size of Reaper is probably limited to whatever size its authors want it to be.

Right now it feels like its authors are experimenting. Building and testing. Maybe Reaper is pure research. We don't know, and that's kind of why we respect it.

Reaper Has Better IoT Security

Unlike many of the devices that it infects, Reaper has an update mechanism. How impressive is that? If it weren't malicious, it might qualify to meet the standards of the new "Internet of Things (IoT) Cybersecurity Improvement Act of 2017" federal requirements. Heck, the authors could even make a distribution out of it and it could become the default remote management platform for IoT.

Is It Malicious?
So far, Reaper hasn't been seen attacking anyone with massive volumetric DDoS attacks. Yes, that's a good thing. At least one of us thinks it might never be seen attacking anyone. If Reaper were to start being used as the ultimate Death Star weapon, that would cheapen its value. It would also result in active takedown campaigns.

Remember how at least two strike-back bots were created to combat Mirai after it attacked Krebs, OVH, and Dyn? Brickerbot actively wiped the filesystems of infected IoT devices (in many cases, turning them into little more than bricks). Hajime was more polite and merely blocked ports and left a cute little note informing the device owner that their device was participating in attacks and please stahp!

If Reaper starts attacking people with DDoS, it will turn from a marvel of thingbot infrastructure engineering into—yawn—another volumetric attack tool. The bot herders would be hunted down by law enforcement (à la the Mirai case3) and the bot would be disassembled.
What Is It Doing?

Right now, Reaper is an object lesson for IoT manufacturers and security researchers. It's like a giant blinking red light in our faces every day warning us that we'd better figure out how to fix IoT security soon.

Is There a Lesson Yet?
As predicted, we will continue to see more thingbots arise as we expect "things" to be the attacker infrastructure of the future. Just because Reaper is the latest, doesn't mean it will be the last. We've added Reaper to the list of botnets that we're monitoring. We suspect that entire existing botnets will get folded into it (whether they wanted to or not).

If Reaper doesn't attack anyone or give away its intentions, it may enter the same mythical space occupied by the Conficker worm of the late 2000s. At its peak, Conficker infected over 10 million Windows computers and caused great concern because it could have done an insane amount of damage. But it was never activated, and it remains a study in bot construction.

The obvious lesson is that the state of IoT security is still incredibly poor, and we need to do a better job of threat modeling the Internet of Things.

~30~