Security Now! #739 - 11-05-19

BlueKeep & DoH

This week on Security Now!

This week we examine a widespread Windows breakage introduced by last month's patch Tuesday. We look at several things Google changed in their just-released Chrome 78, news from the Edge, the status of attacks on Intel chips, a new attack on publicly-exposed QNAP NAS devices, the significant risk of trusting managed service providers, the downside of apps for autos, and worries over Chinese made drones. We then finish by coming back to look at news on two other fronts: The escalating controversy over DNS-over-HTTPS (DoH) and the commencement of the long-awaited BlueKeep vulnerability attacks.

Since October's Patch Tuesday,
Windows will not connect to many websites:

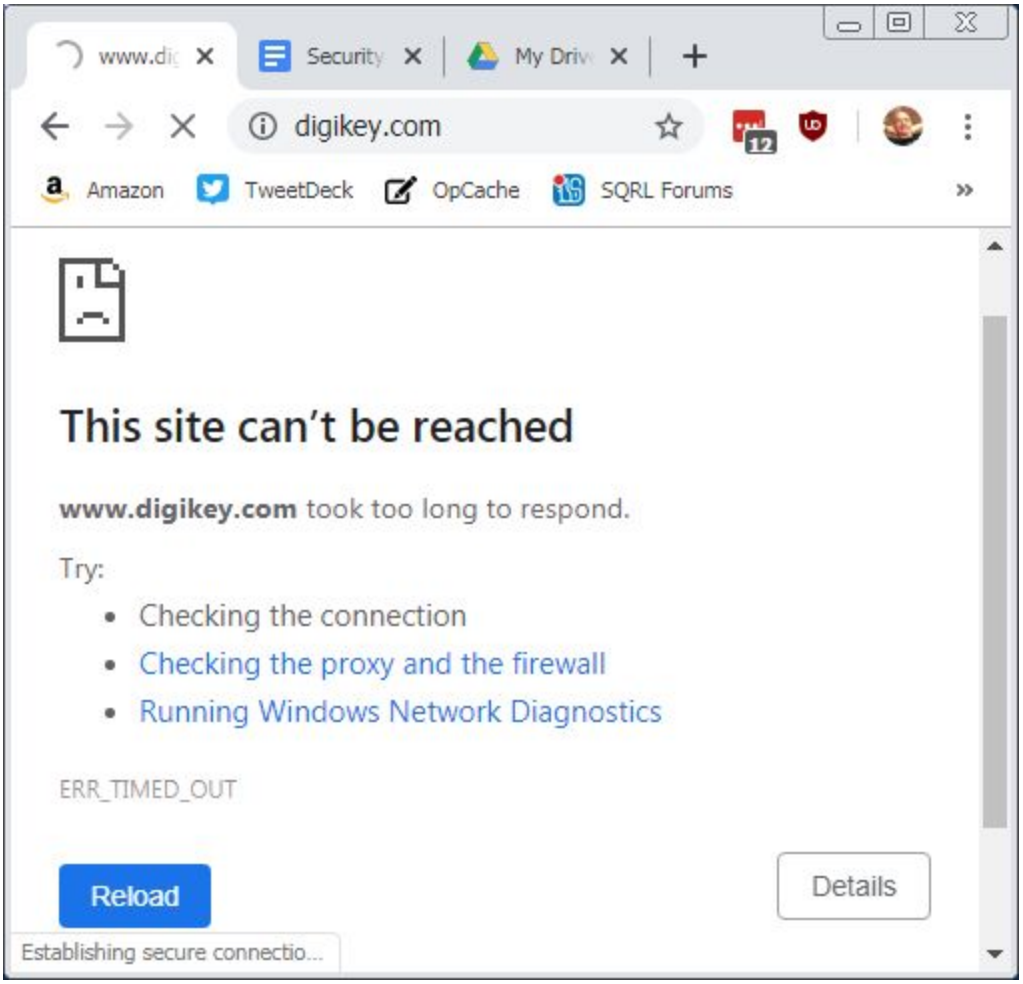

Microsoft Acknowledges.

Security News

Is the world ready for Extended Master Secret (EMS) extensions?
(Apparently not!) A few weeks ago I began having the weirdest trouble with a website that I've been using for years. "DigiKey" is my favorite go-to site for purchasing electronic components online. (Mouser is #2.) As we've recently been discussing, I'll often be using both Firefox and Chrome. I always have an instance of Firefox open, full screen, on my lower-left monitor. It's my static portal to the Internet. And I'll often open an instance of Chrome on my lower-center monitors when I want to do something quick.

So, it's been my habit to pop open a Chrome session when I want to browse around DigiKey for some random part. Except Chrome stopped being able to display DigiKey's site. It would spin and eventually complain that it was unable to obtain a secure connection. This has really been an issue. In the past few weeks I've hit F12 and opened Chrome's developer window to examine the exact error codes being returned.

I've thought that perhaps DigiKey might have somehow misconfigured their site, so I've gone over to Ivan Ristic's SSL Labs and had them scan DigiKey. But DigiKey got top marks from Ivan.

One of the most curious things was that FireFox NEVER had any trouble opening DigiKey. But I wanted to use Chrome for that. And I could not.

This has been bugging me for weeks. And it wasn't just with DigiKey. Under Chrome, random sites all over the Internet were having trouble. Since I learned a long time ago that things that I'm experiencing are likely to be widespread, several times I Googled "Chrome connection errors" thinking that Google must have broken something in Chrome that would be affecting a large number of people.

Then, while catching up on the past week's news for today's podcast, I stumbled upon the reason for ALL of this, and I'm telling everyone who is listening in case you, too, have some favored websites that, for some reason, you haven't been able to get to in Chrome.

So, here it is...

Unbelievable as it is, last month's (October's) Windows Patch Tuesday added support for an enhanced anti-MITM, TLS handshake security feature that immediately BROKE Windows' ability to connect, as a client or as a server, to a significant number of the Internet's websites.

And that perfectly explained why FireFox has continued to work... As we know, FireFox does not use Windows' security stack, it brings along its own. But Chrome DOES run on top of the Windows security protocol stack.

Once I understood what was going on, I immediately disabled this "not yet ready for prime time" protocol handshake feature, rebooted my machine (though that might not have been necessary) and the DigiKey site popped right up under Chrome. Problem solved.

Bleeping Computer started their article on this by writing...

Microsoft has acknowledged a new issue affecting several Windows versions [yes, everything currently being updated, thus all 7 through 10] that could lead to Transport Layer Security (TLS) and Secure Sockets Layer (SSL) connections intermittently failing or getting timed out. [Yeah, or never getting made in the first place.]

This bug is caused by the security-related enforcement for the CVE-2019-1318 TLS spoofing vulnerability, which leads to Windows devices experiencing failures and timeouts during TLS_DHE_* cipher suite negotiation. [DHE are the ephemeral Diffie-Hellman TLS cipher suites]

This happens only when the devices are trying to make TLS connections to devices without support for the Extended Master Secret (EMS) extension.

"Connections between two devices running any supported version of Windows should not have this issue when fully updated," adds Microsoft.

Yeah. Uh huh. Earth to Microsoft: NEWS FLASH! Not everyone, is using your Windows operating systems at each end of the connection on the Internet! Especially not on the server end, where IIS has been steadily losing market share to Nginx, Cloudflare and other providers.

Microsoft's notice of this is headlined: “Transport Layer Security (TLS) connections might intermittently fail or timeout when connecting”

https://support.microsoft.com/en-us/help/4528489/transport-layer-security-tls-connections-might-intermittently-fail-or

Fortunately, a couple of quick Windows registry tweaks can end this trouble by re-disabling Windows new and apparently misguided enforcement of the Extended Master Secret handshake:

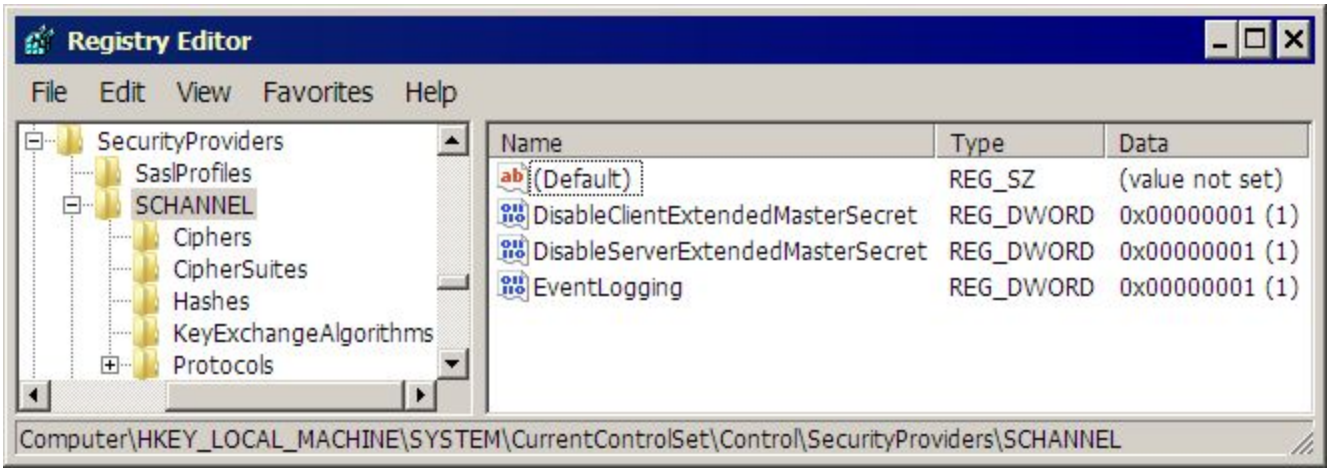

Unbelievable.

And a HUGE and profound THANK YOU to Bleeping Computer for pointing this out. Who knows how long I (and doubtless many others of us) would have been stuck without Chrome working right??

Chrome 78 Disables Code Integrity Check to Mitigate "Aw Snap!" Crashes
While we're talking about fixing things that broke due to overly aggressive security enforcement...

Google backpeddled as fast as they could to officially remove Chrome's enforcement of "Code Integrity Checking" which, as we discussed just last week, suddenly broke Chrome on Windows for anyone who had outdated versions of Symantec's enterprise Endpoint Protection security system. So, Chrome 77 broke it and Chrome 78 unbroke it... by backing off of Code Integrity Checking.

It turned out that three other apps were also causing Chrome to “Aw Snap!” Its users.

After both Symantec and Google provided workarounds for this problem -- which were to update, where possible, to the latest version of SEP or to disable the code integrity protection in the browser -- the “Aw Snap!” crash reports kept on coming!!

The use of Windows Defender Code Integrity Checking prevents binaries that are NOT signed by Microsoft from being loaded by the browser. In an update last week, Google added that "Aw Snap!" crashes were also being triggered on systems running "PC Matic" and Palo Alto Networks' "Traps: security products, as well as the Print Audit Infinite tool which is used for tracking document printing on the LAN.

So Google probably wisely decided to disable CIC in its browser until its compatibility with other software products can be improved. The hope is that this will give developers with products causing this disruption to find a way to fulfill their functions without injecting code into Chrome. Google plans to re-enable the feature in mid-November.

I'll note that since our web browsers are the new front line in the Internet penetration battle, and since this would be a very useful security hardening feature to have enabled in the page renderers of our browsers, if you are a user of any of these known offending add-ons, you would be well served to reach out to the app developer to find out when a fix will be available

And in other Chrome news,
Last Thursday, on Halloween, Google released Chrome 78.0.3904.87, not only to backpeddle on the enforcement of their page renderer's enforcement of code integrity checking, but every bit as urgently to patch a Chrome 0-day which had been discovered by Kaspersky being exploited in the wild!

https://chromereleases.googleblog.com/2019/10/stable-channel-update-for-desktop_31.html

Google posted:

> The stable channel has been updated to 78.0.3904.87 for Windows, Mac, and Linux, which will roll out over the coming days/weeks.
> Note: Access to bug details and links may be kept restricted until a majority of users are updated with a fix. We will also retain restrictions if the bug exists in a third party library that other projects similarly depend on, but haven’t yet fixed.

This update includes 2 security fixes. Below, we highlight fixes that were contributed by external researchers. Please see the Chrome Security Page for more information.

[$7500] [1013868] High CVE-2019-13721: Use-after-free in PDFium.
Reported by banananapenguin on 2019-10-12
[$TBD] [1019226] High CVE-2019-13720: Use-after-free in audio.
Reported by Anton Ivanov & Alexey Kulaev at Kaspersky Labs on 2019-10-29 ***(Note: The SAME DAY as the fix!)***

Google is aware of reports that an exploit for CVE-2019-13720 exists in the wild.

Kaspersky's disclosure offers some wonderfully juicy details:
https://securelist.com/chrome-0-day-exploit-cve-2019-13720-used-in-operation-wizardopium/94866/

Kaspersky Exploit Prevention is a component part of Kaspersky products that has successfully detected a number of zero-day attacks in the past. Recently, it caught a new unknown exploit for Google's Chrome browser. We promptly reported this to the Google Chrome security team. After reviewing of the PoC we provided, Google confirmed there was a zero-day vulnerability and assigned it CVE-2019-13720. Google has released Chrome version 78.0.3904.87 for Windows, Mac, and Linux and we recommend all Chrome users update to this latest version as soon as possible!

Kaspersky endpoint products detect the exploit with the help of the exploit prevention component. The verdict for this attack is Exploit.Win32.Generic.

We are calling these attacks Operation WizardOpium. So far, we have been unable to establish a definitive link with any known threat actors. There are certain very weak code similarities with Lazarus attacks, although these could very well be a false flag. The profile of the targeted website is more in line with earlier DarkHotel attacks that have recently deployed similar false flag attacks.

More details about CVE-2019-13720 and recent DarkHotel false flag attacks are available to customers of Kaspersky Intelligence Reporting. For more information, contact: intelreports@kaspersky.com.

Technical details
The attack leverages a waterhole-style injection on a Korean-language news portal. A malicious JavaScript code was inserted in the main page, which in turn, loads a profiling script from a remote site.

The main index page hosted a small JavaScript tag that loaded a remote script from hxxp://code.jquery.cdn.behindcorona[.]com/.

The script then loads another script named .charlie.XXXXXXXX.js. This JavaScript checks if the victim's system can be infected by performing a comparison with the browser's user agent, which should run on a 64-bit version of Windows and not be a WOW64 process; it also tries to

get the browser's name and version. The vulnerability tries to exploit the bug in Google Chrome browser and the script checks if the version is greater or equal to 65 (current Chrome version is 78):

If the browser version checks out, the script starts performing a number of AJAX requests to the attacker's controlled server (behindcorona[.]com) where a path name points to the argument that is passed to the script (xxxxxxx.php). The first request is necessary to obtain some important information for further use. This information includes several hex-encoded strings that tell the script how many chunks of the actual exploit code should be downloaded from the server, as well as a URL to the image file that embeds a key for the final payload and RC4 key to decrypt these chunks of the exploit's code.

After downloading all the chunks, the RC4 script decrypts and concatenates all the parts together, which gives the attacker a new JavaScript code containing the full browser exploit. To decrypt the parts, the previously retrieved RC4 key is used.

The analysis we have provided here is deliberately brief due to vulnerability disclosure principles. The exploit used a race condition bug between two threads due to missing proper synchronization between them. It gives an attacker an a Use-After-Free (UaF) condition that is very dangerous because it can lead to code execution scenarios, which is exactly what happens in our case.

The exploit first tries to trigger UaF to perform an information leak about important 64-bit addresses (as a pointer). This results in a few things: 1) if an address is leaked successfully, it means the exploit is working correctly; 2) a leaked address is used to know where the heap/stack is located and that defeats the address space layout randomization (ASLR) technique; 3) a few other useful pointers for further exploitation could be located by searching near this address.

After that it tries to create a bunch of large objects using a recursive function. This is done to make some deterministic heap layout, which is important for a successful exploitation. At the same time, it attempts to utilize a heap spraying technique that aims to reuse the same pointer that was freed earlier in the UaF part. This trick could be used to cause confusion and give the attacker the ability to operate on two different objects (from a JavaScript code perspective), though in reality they are located in the same memory region.

The exploit attempts to perform numerous operations to allocate/free memory along with other techniques that eventually give the attackers an arbitrary read/write primitive. This is used to craft a special object that can be used with WebAssembly and FileReader together to perform code execution for the embedded shellcode payload.

After decryption, the malware module is dropped as updata.exe to disk and executed. For persistence the malware installs tasks in Windows Task Scheduler.

So, in short, from the user's standpoint... With Chrome before 78, you visit a website and without you knowing or clicking or doing anything, your Chrome browser downloads, installs and runs a native malware application on your Windows PC.

And... We are at the first Microsoft Edge Stable Release Candidate!
A preview build for the new Chromium-based Microsoft Edge Stable version has been released on the Edge Insider site. This is the preview of the Edge browser that will be made widely available in January.

Monday, kicking off the start of Microsoft's five-day Ignite 2019 conference, Microsoft's Yusuf Mehdi, Corporate Vice President, Modern Life, Search and Devices Group, stated that Microsoft Edge will become generally available and fully released, on January 15th, 2020.

With this preview release, Microsoft Edge Stable is currently at version 79.0.309.7, which puts it one major version ahead of Google's Chrome current stable version, which, as we just noted in the previous news, has just moved to 78.

From an academic curiosity standpoint it's going to be very interesting to see how the world settles out with a version of Edge that's internally identical to Chrome. Will Edge take some of Chrome's market share? Will anyone know or care what's under the hood of their browser? What has been moving people away from Edge in the first place? I have no idea. But it's going to be interesting to see how this develops. In general, I applaud this move by Microsoft. It must have been so difficult to discard all that work that was done on their own new browser. But this was the sanest choice.

Microarchitectural Data Sampling Vulnerabilities
That's the fancy roll-up catch-all classification name that groups together all of the side-channel attacks on Intel and, to a lesser degree AMD and other processors which started coming to light right from the start of 2018, nearly two years ago.

So we first had Meltdown and Spectre, they were followed by a few variants. Then we were hit by PortSmash, ZombieLoad, RIDL, and Fallout.

We're nearly two years downstream of all of that and one glaring aspect remains true: Despite there never having been any sign of any of these theoretical processor failings being leveraged into a practical attack, we have nevertheless sacrificed around 20 percent of our previous system performance in order to prevent these attacks that never materialized.

Of course, we can't prove a negative. So there's no way to know whether, if the industry hadn't responded as it did, we would have eventually discovered these problems in the wild. We'll likely never know.

Hopefully, Intel has learned a lot from all of this and, although their processor design pipeline is very long and deep, so that it will take some time for a fully post-Specter and Meltdown new processor design to emerge, we may eventually have Microarchitectural Data Sampling hardened processors for our future machines. At which time we'll likely be able to recover that lost 20% of our system's performance.

Trouble for QNAP NAS devices exposed to the Internet

I know from many instances of feedback I've seen after our podcast about my search for the right file synchronization solutions, that a great many of our listeners are happy users of the QNAP brand of network attached storage devices. So I wanted to make sure that they and everyone were aware of the fact that many thousands of QNAP NAS devices have been infected with the QSnatch malware. Over 7,000 infections have been reported in Germany alone and the malware is still spreading. Thousands more are believed to be infected worldwide, in what appears to be an ongoing outbreak.

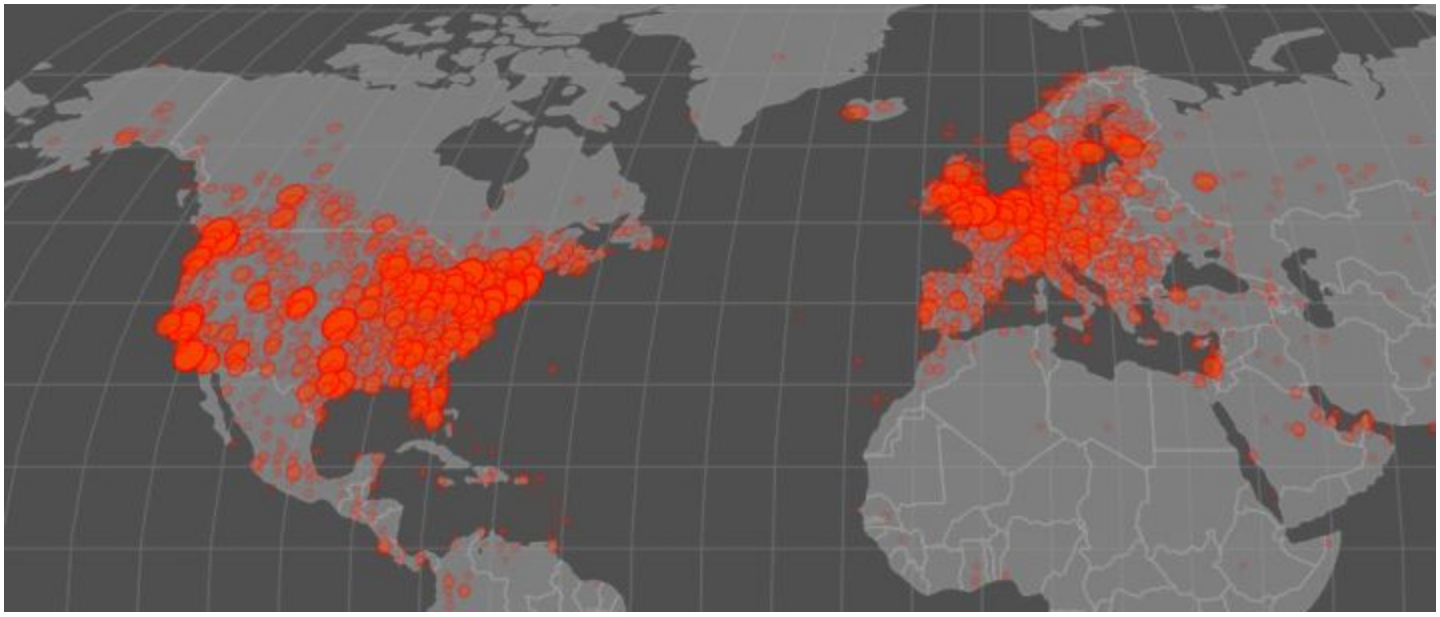

Information on how QSnatch works is still scant. The only report comes from the National Cyber Security Centre of Finland (NCSC-FI), which was the first cybersecurity organization to spot the malware last week.

NCSC-FI members have not yet discovered how this new threat spreads and infects QNAP NAS systems; however, once it gains access to a device, QSnatch burrows into the firmware to gain reboot persistence.

An analysis of the malware's code revealed the following capabilities:

- Modify OS timed jobs and scripts (cronjob, init scripts)
- Prevent future firmware updates by overwriting update source URLs
- Prevents the native QNAP MalwareRemover App from running
- Extracts and steals usernames and passwords for all NAS users

These features describe the malware's capabilities but don't reveal its end-goal. It is unclear if QSnatch was developed to carry out DDoS attacks, to perform hidden cryptocurrency mining, or just as a way to backdoor QNAP devices to sensitive steal files or host malware payloads for future operations.

One theory is that the QSnatch operators are currently in the phase where they're building their botnet, and will deploy other modules in the future. NCSC-FI analysts confirmed that QSnatch

has the ability to connect to a remote command-and-control, download, and then run other modules.

For the time being, the only confirmed method of removing QSnatch has been performing a full factory reset of the NAS device. After performing a factory reset, users should install the latest QNAP NAS firmware available. Last Friday, November 1st, (2019) QNAP released a firmware update which incorporates specific QSnatch protections. Anyone with an Internet-exposed QNAP NAS will want to be current.

Other advice shared by NCSC-FI analysts on dealing with the aftermath of a QSnatch infection include:

- Change all passwords for all accounts on the device
- Remove unknown user accounts from the device
- Make sure the device firmware is up-to-date and all of the applications are also updated
- Remove unknown or unused applications from the device
- Install QNAP MalwareRemover application via the App Center functionality
- Set an access control list for the device (Control panel -> Security -> Security level)

QSnatch is the fourth malware strain spotted this year that has targeted NAS devices, following in the footsteps of a ransomware strain that impacted Synology devices, and the eCh0raix and Muhstik ransomware strains that previously impacted QNAP devices.

MSP's -- Managed Service Providers -- are a major vector for ransomware delivery
This year, at least 13 managed service providers, who have each been positively identified, were used as the means to push ransomware out to their client companies. The MSP functions as a large access hub. Once hackers compromise an MSP's network, they use its remote access tools to deploy ransomware to hundreds of companies and thousands of computers.
https://www.armor.com/reports/new-msps-compromised-reports-armor/

The global cloud security solutions provider whom we have referred to in the past, Armor, has just published a report identifying 6 new Managed Service Providers (MSPs) and Cloud-Based Service Providers that have been compromised by ransomware, which bringings the total number of publicly identified MSPs and Cloud-Based Service Provider victims in 2019 to 13.

I'm not going to go though an enumeration of this. It's all there for anyone who does want a list and additional information. It's the concept and the threat it represents that I wanted to be sure our listeners understand and appreciate.

Chris Hinkley, Head of Armor’s Threat Resistance Unit (TRU) research team, said: “This uptick in successful ransomware attacks against MSPs and Cloud-Based Service Providers is a harsh reminder that organizations need to ensure that the third-party vendors they do business with are as equally protected against current and emerging cyber threats, as they are. This is especially true, because as we have seen, a successful ransomware attack against a MSP/Cloud-Based Service Provider can be debilitating to their customers, as well as to their own company, as the attack can quickly shut down key systems which the customers depend on to run their organization.”

”And of course, a ransomware attack against an MSP can be fatal, putting a MSP out of business, which appears to have happened with "PM Consultants", the Oregon-based IT consulting and IT support provider to dental practices. After they were hit by ransomware in early July they subsequently shut down their business down later that month, citing that they were doing so in part due to the ‘devastating event‘.”

So, our takeaway here is that any enterprise who is contracting for the services of any outside provider should give serious consideration to the sort of access that the provider truly needs into their network. It's entirely natural for any provider to have full trust in their own capabilities and security, and so to in good conscience ask for full and unfettered access to their client's networks. But any highly responsible IT manager who was on the receiving end of such an arrangement should give some serious consideration to what sort of access the outside party needs, understanding that something entirely inadvertent on the other end could happen.

It's not possible to offer anything more concrete without knowing exactly what the relationship is and what form of access is really needed. But we are seeing many instances where, whatever level of access IS being provided, it's like much more than the MSP's victim client WISH they had provided after the fact.

Five months after returning rental car, man still has remote control
Man can still track vehicle, lock and unlock it, and start and stop its engine.

When Masamba Sinclair rented a Ford Expedition from Enterprise Rent-a-Car last May, he was excited to connect it to FordPass. The app allows drivers to use their phones to remotely start and stop the engine, lock and unlock the doors, and track the vehicle's precise location.

[What could POSSIBLY go wrong???]

34 year old Masamba said: "I enjoyed it and logged into FordPass to be able to access vehicle features from my phone such as locking, unlocking, and starting the engine. I liked the idea of it more than I found it useful. The UI looks good and works well, though."

Today, Sinclair's opinion of mobile apps in rental cars is decidedly less favorable. That's because, five months after he returned the vehicle on May 31, his app continues to have control over the vehicle. Despite multiple other people renting the SUV in the intervening months, FordPass still allows Sinclair to track the location of the vehicle, lock and unlock it, and start or stop its engine. Sinclair has brought the matter to Ford's attention, both through its website and multiple times on Twitter. So far, Ford has done nothing to end his access.

He said: “All it took for me to initially connect was to download the app and enter the VIN number, then confirming connectivity through the infotainment system. There MIGHT be a way to disassociate my phone from the car itself, but that hasn't happened yet, and it's crazy to put the onus on renters to have to do that. I have had no problems at all and have even unlocked the doors and started the engine when I could see that the vehicle was in the Missoula airport rental car parking lot.”

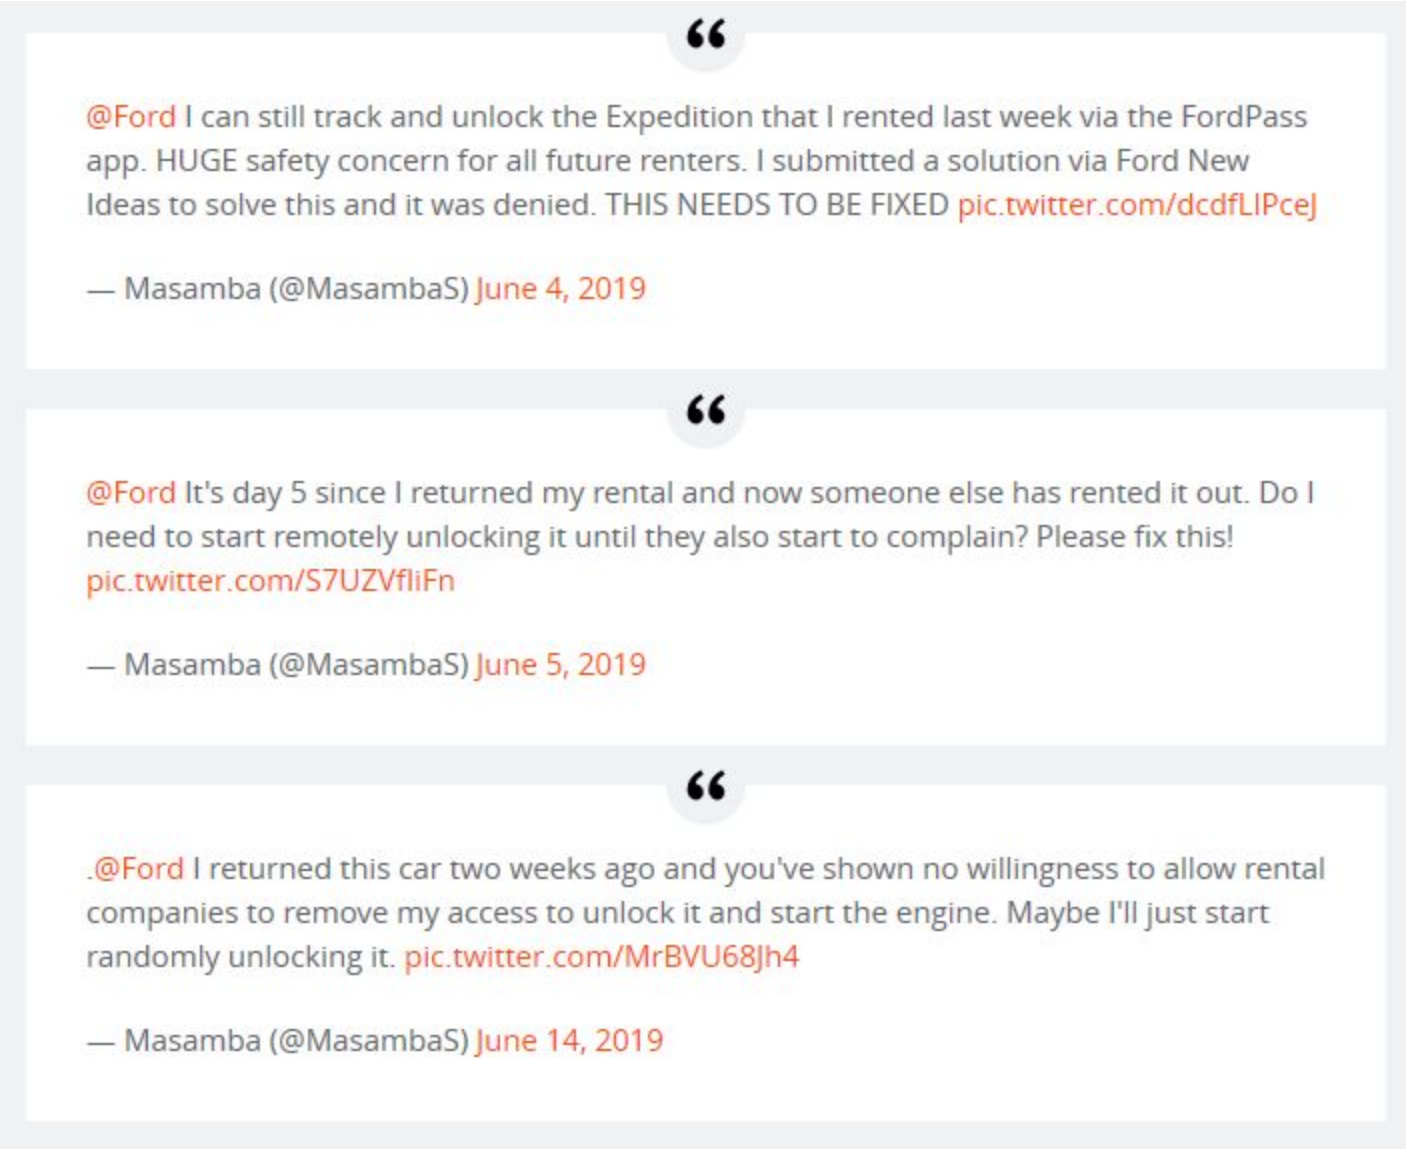

FordPass is offered by the Ford Motor Company and is available for both iOS and Android devices. It is one of several apps for connecting to Ford vehicles. The less-than-intuitive means for unpairing a vehicle and phone—not to mention the difficulty in knowing a device remains connected—represent a serious security and privacy risk, not just to renters, but to people buying a vehicle second hand.

While Ford said infotainment screens will indicate when a device is paired, it's obvious that multiple Enterprise employees and renters have continued to miss the warning. Even now, after the reporter of this discussed the problem with both Enterprise and Ford representatives, Sinclair's access still hasn't been revoked.

Sinclair said: "I have been opening the app and tracking the vehicle almost every day to see if my access is still there, and sure enough, I can see exactly where my old rental, affectionately named "The Beast," is at any given moment. This means that I can not only find this rental car whenever I want, but I can also unlock the doors and help myself to anything inside."

Since proximity to the infotainment system was required to complete the initial pairing authorization and authentication, and since occasional proximity would be an expected

characteristic for any actual car owner, it would seem to me that a useful security tradeoff would be to have a device be forgotten if it hasn't been within physical proximity of the vehicle for, say, a week or two.

Chinese-made Drones in the US are being grounded

I have no way of independently assessing the danger that foreign-made drones could pose to US domestic security. But they are flying everywhere and they are essentially flying video cameras. And they are connecting to the Internet and they are sending data back to China.

So what's in the news is that, pending a security review of Chinese-made drones, the US Department of the Interior (DOI) announced last Wednesday that it's grounding all Chinese-made drones and drone with Chinese-made parts as it reviews its drone program. And the Department of the Interior does have quite a program.

But because it's simply not practical, that decision does not apply to drones "currently being utilized for emergency purposes, such as fighting wildfires, search and rescue, and dealing with natural disasters that may threaten life or property.

In recent years, the DOI has enthusiastically embraced drones, publicizing the wide variety of ways it has deployed them. In addition to being deployed for emergency rescues and disaster monitoring, drones are used in more expansive, long-term projects such as geological surveys and wildlife population monitoring.

According to a 2018 report about its use of drones, the department owned 531 drones as of last year, and conducted more than 10,000 flights across 42 states and territories. Other more recent reports have updated that number to 800 operating drones. The report did not specify what percentage of those drones were Chinese-made. However, the most popular and capable Drones are made by the Chinese company DJI. And, in fact, two years ago the US Army discontinued the use of drones produced by DJI, the world's biggest manufacturer of drones, over the risk of vulnerabilities. And last May, the Department of Homeland Security warned companies that Chinese-made drones could potentially transmit sensitive footage or data to third parties. And last month, a bipartisan group of lawmakers introduced legislation that would bar all federal agencies from operating drones manufactured or assembled in China.

Because the Interior Department uses drones to survey a variety of critical infrastructure, including mines and dams, as well as to study rapid response situations and emergency routes, the information they collect has at least some potential for abuse.

To date, scant public evidence exists that Chinese drones have been involved in large-scale cyberespionage, or that they have back doors built in that would allow them to be exploited for surveillance. For their part, DJI has pushed back and argued that opposition to its products is motivated primarily by political hostility toward Chinese companies... but we all know that, just as with hiding itty-bitty malicious chips on motherboards... even if it isn't being done, it could be done. So I suppose I'm glad that at least a sober awareness of the risks exists and is being taken seriously.

DJI has proposed assembling more of its drones in the United States, as well as offering a

version of its devices, called the Government Edition, with certain safeguards that would protect information captured would and prevent it from being transmitted wirelessly.

What may develop is a two-tier system where governments, which can afford to pay for more costly, audited and known-clean drone will have that option, but wedding photographers can save their money, since recording Jack's and Jill's nuptials from a height is unlikely to be of great interest to foreign spies. And DJI could continue to make their lovely hardware, but allow an open-source community to add the software.

DoH

The DoH Controversy
https://www.zdnet.com/article/mozilla-cloudflare-doesnt-pay-us-for-any-doh-traffic/

DoH remains controversial -- not so much with end-users who appreciate the idea that their ISPs and others who might wish to snoop on and possibly filter and block their DNS lookups will be much less able to do so -- but within the industry where those who are being blinded to DNS lookups are crying foul and running around stirring up a bunch of FUD - Fear, uncertainty and doubt.

For example week before last, Motherboard carried a story with the headline: "Comcast Is Lobbying Against Encryption That Could Prevent it From Learning Your Browsing History"

And the sub-head reads: "Motherboard has obtained a leaked presentation internet service providers are using to try and lobby lawmakers against a form of encrypted browsing data."

First, let's remember as we have previously covered here, Mozilla has indicated that it plans to have Firefox route its DoH queries to Cloudflare's servers. And in response to the flack it's been getting, Mozilla has been explicit that <quote> "no money is being exchanged to route DNS requests to Cloudflare" as part of the DNS-over-HTTPS (DoH) feature that is currently being gradually enabled for Firefox users in the US. Detractors have been saying that by using Cloudflare as the default DoH resolver for Firefox, Mozilla will help centralize a large chunk of DNS traffic on Cloudflare's service.

And let's also remember, as we've also described previously, that Google is planning to take a more dynamic posture. It will have Chrome test the user's currently-configured DNS and IF that service also supports DoH, Chrome will switch to using it. This is a slick solution since Chrome users may have already manually switched their non-encrypted DNS to some other provider who offers superior value-added services. So, if that's the case, why not allow it to also be encrypted? So let's keep this in mind when we look at what COMCAST is alleging to lawmakers...

Motherboard:

> The plan, which Google intends to implement soon, would enforce the encryption of DNS data made using Chrome, meaning the sites you visit. Privacy activists have praised Google's move. But ISPs are pushing back as part of a wider lobbying effort against encrypted DNS, according to the presentation. Technologists and activists say this encryption would make it harder for ISPs to leverage data for things such as targeted advertising, as well as block some forms of censorship by authoritarian regimes.
>
> Mozilla, which makes Firefox, is also planning a version of this encryption.
>
> Marshall Erwin, senior director of trust and safety at Mozilla, told Motherboard in a phone call after reviewing sections of the COMCAST slide deck that: "The slides overall are extremely misleading and inaccurate, and frankly I would be somewhat embarrassed if my team had provided that slide deck to policy makers." He added that: "We are trying to essentially shift the power to collect and monetize peoples' data away from ISPs and providing users with control and a set of default protections."

Comcast:

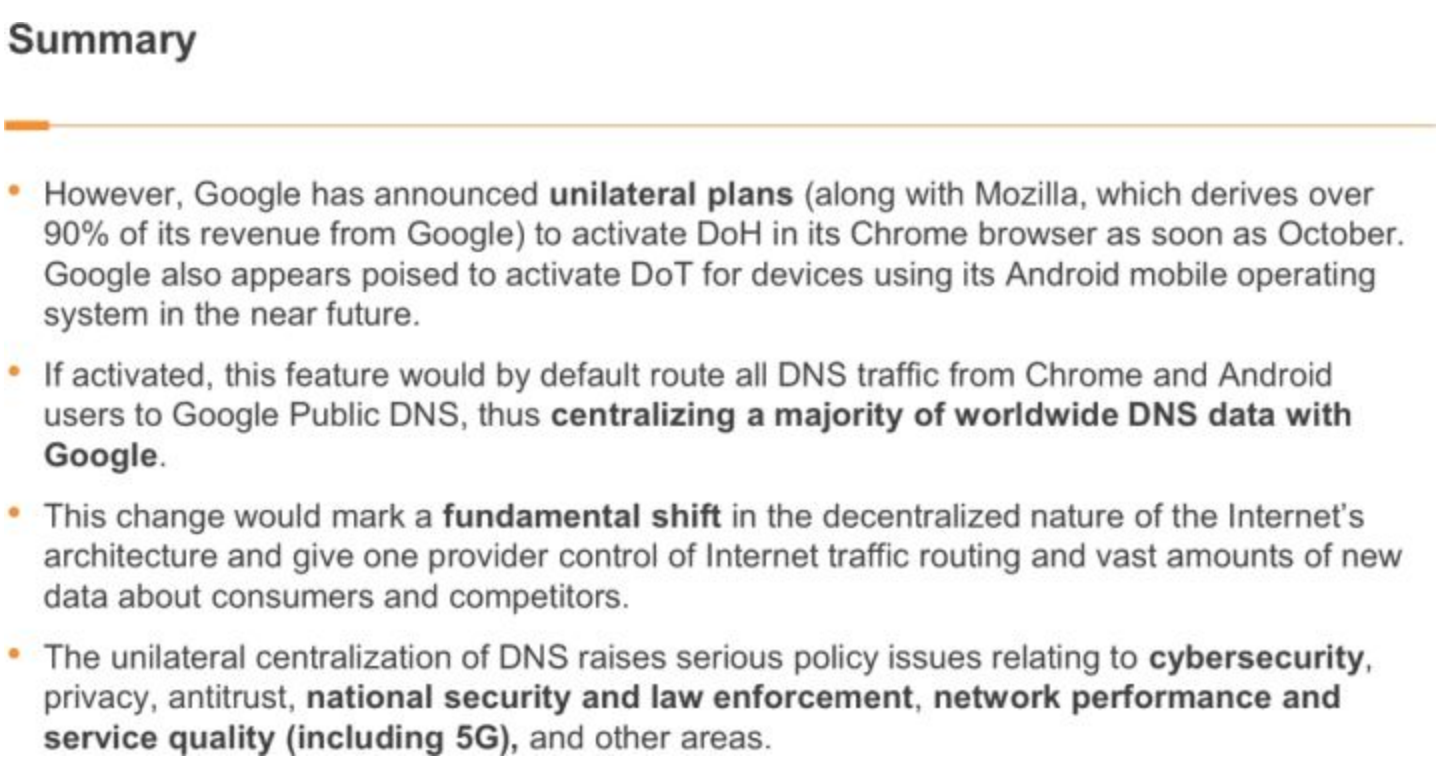

In the presentation, Comcast paints this type of encryption as something that will fundamentally change the internet and will centralize power under Google.

Comcast wrote in its presentation: “The unilateral centralization of DNS raises serious policy issues relating to cybersecurity, privacy, antitrust, national security and law enforcement, network performance and service quality (including 5G), and other areas. Congress should demand that Google pause and answer key questions. Why is Google in such a rush?”

The entire presentation slide deck is embeded in the Motherboard article, which anyone listening can view if they're interested, and I have a link to the PDF in the show notes:

https://assets.documentcloud.org/documents/6509454/ISP-DoH-Lobbying-Slide-Deck.pdf

But to give everyone a feel for the deck's content, here are the headings of the pages. And as I read these, it's worth keeping in mind that back in 2017 ISPs lobbied Congress to make it possible to sell your browsing data without your consent.

The presentation is titled: "Google's Proposed Public DNS Plans"

1. Google's unilateral imposition of default, CENTRALIZED DNS encryption will harm key components of the Internet.
2. Google's plans will cause radical disruption.
3. Significant cybersecurity and national security risks.
4. Privacy risks increase as a single firm amasses more consumer data.
5. Antitrust and competition concerns.
6. Law enforcement efforts may be compromised.
7. CDN localization will likely suffer and backbone costs will rise.
8. Wireless performance, including 5G, will be undermined.
9. Parental controls & content filtering may be disabled.
10. ISPs and other enterprise services may be disrupted or broken.
11. Congress should demand the Google pause and answer key questions.

We've **REALLY** got to ask ourselves why COMCAST apparently places so much value upon their access to their customers' DNS data.

BlueKeep

So, it took a long time, but the BlueKeep-based attacks have finally started, and what we predicted on this podcast has finally happened. The so-called "BlueKeep" vulnerability in Windows Remote Desktop is being exploited, and not by a worm.

ZDNet's headline: "BlueKeep attacks are happening, but it's not a worm. Hackers are using BlueKeep to break into Windows systems and install a cryptocurrency miner."

The Hacker News: "First Cyber Attack 'Mass Exploiting' BlueKeep RDP Flaw Spotted in the Wild"

Threatpost: "BlueKeep Attacks Have Arrived, Are Initially Underwhelming"

Bleeping Computer: "Windows BlueKeep RDP Attacks Are Here, Infecting with Miners"

Over the weekend and across the world, security researcher's honeypots monitoring port 3389 began lighting up as attempts were being made to leverage the BlueKeep vulnerability for the purpose of running cryptocurrency miners.

On Saturday, Kevin Beaumont noticed that multiple honeypots in his EternalPot RDP honeypot network started to crash and reboot. His "Pots" have been active for almost half a year, so they would have been scanned and cataloged by anyone who was planning an attack. And during all that time, this is the first time they came down.

The first details about BlueKeep being the cause of these events came from Marcus Hutchins (aka MalwareTech), who investigated the crash dumps from Beaumont's machines. He said that he “found BlueKeep artifacts in memory and shellcode to drop a Monero Miner.”

According to his analysis, an initial payload runs an encoded PowerShell command that downloads a second PowerShell script, also encoded. Marcus says that the final payload is a cryptocurrency miner, likely for Monero, currently detected by 25 out of 68 antivirus engines at VirtusTotal.

Marcus said that the malware may not be a worm but that it is mass-exploiting the BlueKeep bug. This indicates that the cybercriminals are using a BlueKeep scanner to find vulnerable systems exposed on the web and drop the cryptocurrency miner on them. And in a later update Marcus said that analysis of its network traffic does not indicate self-propagation.

So the server doing the exploitation obtained its target IP addresses from a predefined list. This is what we expected since the exploit is straightforward once its details have been worked out, and adding worm code just adds additional unnecessary complexity without returning any real value. Once upon a time, when the Internet was a much more quiet place, a worm would have been launched as a means of obscuring the original attacker’s origin. But on today’s Internet, where other’s machines are easily compromised, hiding it far less necessary.

The first public BlueKeep exploit was added to Metasploit two months ago in September and Marcus' analysis has confirmed that the same code which is present in the open source Metasploit module is also present in the malware.

So it appears likely that whoever is behind these attacks is using publicly available code resources and did not develop a reliable, wormable threat, as Kevin Beaumont's honeypot crashes suggests. Kevin has recently noted that there are presently more than 724,000 systems worldwide which are susceptible to BlueKeep exploitation.