Security Now! #723 - 07-16-19

Encrypting DNS

This week on Security Now!

This week we cover a few bullet points from last last Tuesday's monthly Windows patches as well as some annoyance that the patches caused for a Windows 7 users, we track some interesting ongoing Ransomware news, we look at the mixed blessing of fining companies for self-reporting breaches, we check out a survey of enterprise malware headaches, update on some Mozilla/ Firefox news, examine yet another (and kinda obvious) way of exfiltrating information from a PC. We address a bit of errata, some miscellany and closing the loop with our listeners, then we conclude with a closer look at all the progress that's been occurring quietly with DNS Encryption.

Security News

Patch Tuesday Update

Two zero days and 15 critical flaws fixed in July's Patch Tuesday
https://nakedsecurity.sophos.com/2019/07/10/two-zero-days-and-15-critical-flaws-fixed-in-julys-patch-tuesday/

16 critical vulnerabilities, some being exploited, fixed in July, 2019 Windows updates

July's security updates address 77 vulnerabilities that affect Windows and a range of software that runs on Windows, mainly Internet Explorer, DirectX and Windows' graphical subsystem.

Of those 77 vulnerabilities, Microsoft rates 16 of those as critical, 60 as important and 1 as moderate.

Most of the critical vulnerabilities allow attackers to execute remote code on the user's system, and 19 of the important vulnerabilities can be used for local elevation of privilege -- which, as have seen, is really not much less threat even though it sounds much more benign.

6 of the critical vulnerabilities are for IE itself and 5 are for Chakra, which as we know is the JavaScript engine in both original (non-Chromium) Edge and IE. In addition, the components that all have one remote code execution vulnerability are: Windows' DHCP server, the Azure DevOps Server, the .NET Framework and the GDI+ API.

And that one "moderate" problem resolved was an authentication bypass for applications using the Windows Communication Foundation and the Windows Identity Foundation API.

And... It appears that two of the elevation of privilege vulnerabilities were 0-day vulnerabilities that were being actively exploited by Russian attackers. As I said, EoP (elevation of privilege) is extremely handy.

The two 0-days which will be under active exploitation until machines are updated are:

Win32k Elevation of Privilege Vulnerability / CVE-2019-1132
The Win32k driver on Windows 7 32bit can be abused to get a NULL pointer dereference. An attacker with remote code execution could use this vulnerability to achieve local elevation of privilege. This has been exploited in the wild.

Microsoft splwow64 Elevation of Privilege Vulnerability / CVE-2019-0880
There is pointer dereference issue in the printer driver for 32-bit processes that could be used to escape the sandbox of Internet Explorer Enhanced Protected Mode (EPM). After achieving remote code execution using a vulnerability like the ones that follow, affecting the Edge or Internet Explorer web browsers, an attacker could exploit this vulnerability to create a new process of medium integrity level. This has been exploited in the wild.

As for how the attacker gets in to elevate their privilege, there were 6 browser memory corruption vulnerabilities in IE and another 5 in the JavaScript Chakra engine.

Internet Explorer and Edge suffer from several memory corruption vulnerabilities, such as type confusion, out-of-bounds write, and use-after-free.

If an attacker were to cause a victim to visit a malicious website, they could execute remote code in the context of the web browser. To gain full control of the machine is more difficult: the attacker would need to escape the sandbox using a vulnerability such as the previously mentioned CVE-2019-0880, and then perform a local privilege elevation that delivers full administrative access.

As an aside, Adobe did not synchronize the release of their own patches this month as they often do.

When is a security update not a security update?
On a related note was the news that some stalwart Windows 7 users got very worked up and annoyed (to put it mildly) to receive a non-security update which added new Windows telemetry to their Win7 machines, despite the fact that it was labeled as a security-only monthly patch.

Recall that back in 2016 Microsoft simplified its ongoing patching of Windows versions by offering Windows 7 and 8.1 users two types of update: Either the 'Monthly Rollup' containing both security and non-security patches (i.e bugs and reliability), the second a security-only update option which only repaired that month's security flaws. So it turns out that Microsoft's July 9, 2019 KB4507456 (Security-only update) actually contained more than only security. Something called the "Compatibility Appraiser" tool was slipped into that update.

Our friend Woody Leonard, writing in his "Woody on Windows" column in ComputerWorld, posted on July 11th under the title: "New Windows 7 'security-only' update installs telemetry/snooping, uh, feature"

https://www.computerworld.com/article/3408496/new-windows-7-security-only-update-installs-telemetrysnooping-uh-feature.html

His piece's subhead reads: "Three years ago, Microsoft promised to keep Win7 and 8.1 updated with two tracks of patches - Monthly Rollups that include everything and "security-only" patches that are supposed to be limited to security fixes. Guess what just happened. "

Woody's article has a ton of good information for those who want to know more. In it he cites a Dr. Vess Bontchev who tweets as @VessOnSecurity:

"I have officially stopped updating my Win7 machine. I no longer trust Microsoft's updating process. I'll protect it from any existing and future vulnerabilities with my other defenses, as well as I can. F**k you, @microsoft" (Woody politely left off the F-U @Microsoft.)

My feeling is, that all we can be is informed. So that's what we do on this podcast. (For example, I received a huge amount of positive feedback from my observation that our SOHO routers could be configured to forward all DNS queries over HTTP or TLS.) So that's why we're here -- to stay informed. I choose to use Windows 7, which I do with my eyes wide open. The job Microsoft is doing is impossible. I don't want that job. No one wants it. And given the messy legacy of

Windows code, its barely Windows-literate user-base, and the incredibly and increasingly hostile environment from which Windows attempts to protect its users, they're doing an amazing job. Sure, they're being paid handsomely for that job. But they're doing it.

I feel as though I should take this moment to talk a bit about Windows 7, Windows 10 and me... because security updates **will** stop flowing to Windows 7 six months from now, next February -- unless, that is, Microsoft changes their minds again and pushes that deadline back still further, because they continue failing to force the world to use a version of Windows it doesn't want.

As we know, at the beginning of the year, Windows 10 finally changed places with Windows 7 as first place desktop OS. But today, six months later, they are still neck and neck with Win10 sitting at 40.61% with Windows 7 at 38.06%. So they are still near parity.

As our long-time Security Now! Listeners know, I won't be moving my main Workstations to Win10 once Win7 stops being supported. I have Win10 on laptops for testing. When I go out to present SQRL to a group, that's Win10 running on that laptop. And we're Skyping over Windows 10. But I know from long, loving and trouble-free experience with XP, that this Win7 machine I'm sitting in front of, with its five high-resolution screens, will continue to happily purr away for many many years even without constant nursing from mama.

Windows Defender has never found anything on any of my machines other than false-positive annoyances with my own code that it doesn't know about, or old well-marked viruses in eMail archives from those good old days. But I'll miss having Windows Defender officially integrated into Windows 7 and watching my back. It's comforting. So I'll be giving up some comfort. But I expect our browser defenses to continue to improve, and all non-Microsoft browsers will continue running on and supporting Windows 7 for many more years. So I'll be okay. Remember that those 11 web browser vulnerabilities in IE and Edge were the way bad guys were getting in. So I will continue to **not** use IE. I'll be sticking with Firefox for the foreseeable future.

And, after all, my backups have backups and I keep rolling, off-machine, incremental file backups of every project I'm working on as well as monthly static deep-freeze snapshot images. So I'm pretty well protected. But I'm not your average user, and for what it's worth, neither are the listeners of this podcast. I love so many of the apps I have running on Windows -- I'm extraordinarily happy with them. There are many that I have moved forward from machine to machine through the years. The move away from XP, and the loss of native support for 16-bit apps, was traumatic. But it finally had to be done because even the web browsers had finally started refusing to update. So the death of that machine was a mixed blessing. Consequently, I fully expect to remain where I am using Windows 7 until something -- other than a lack of monthly nursing -- forces me to move.

But just so I'm clear, even with all that said, I'm **NOT** suggesting that anyone else follow my example. I'm really not. My use of Windows is **boring** compared with most others. And it's often the case that "boring" is "secure." I don't use my machine for entertainment or gaming, or watching YouTube videos or, out of boredom, randomly clicking links to see what's out there. I'm really not very interested in most of what is out there. So while my main Win7 workstations are not technically "Air-Gapped" from the Internet, they are "Steve-Gapped" -- which exposes them to many fewer threats than the typical Windows PC.

La Porte County, Michigan struck by Ryuk...
The Michigan City News Dispatch reported last Tuesday, July 9th: "Malware attack on county computers. La Porte County website, government email servers out of operation"

https://www.thenewsdispatch.com/news/article_d9809e48-7e8d-52d5-9d08-5d6c1adab2a2.html

La PORTE – (Paraphrasing and trimmed) All La Porte County government emails, and the county website, remained out of commission late Tuesday following a malware virus attack that affected the system on Saturday morning.

La Porte County Board of Commissioners President Dr. Vidya Kora said Sunday evening, the system will be inoperable as authorities respond to a "malicious malware attack that has disabled our computer and email systems."

County Attorney Shaw Friedman confirmed [on] Tuesday that county government computers were "impacted by a sophisticated ransomware virus" early Saturday morning. He said: "Fortunately, our IT team reacted quickly and shut down much of the system, even though it was a weekend. Less than 7 percent of our laptops have been infected, however, it did hit our two domain controllers, which means no server can access network services."

An insurance policy taken out last year will help the county recover, Kora said, "Fortunately, our county liability agent of record, John Jones, last year recommended a cybersecurity insurance policy which the county commissioners authorized from Travelers Insurance. We informed Travelers Insurance late Saturday of the malware attack and they immediately referred us to the Wayne, Pennsylvania, incident-response law firm of Mullen Coughlin LLC that specializes in responses to such cyber-attacks and coordinates system repairs and protection of our computers from further such virus infections."

Friedman said: "The forensic investigation firm has been retained to determine the nature and scope of the incident, including how the county could have been infected. We're developing a game plan to respond to the attack and come up with an approach to repair our systems and protect them from further damage. The county's IT Department has been working long hours to try and get things operational, including spending Sunday to ensure that the Courts and Prosecutor's office remained functional."

"This particular ransomware variant – known as RYUK – is especially insidious as it seeks to delete or encrypt system backups. We are exhausting all possibilities, including tapping the FBI cybersecurity unit and reviewing all 'workarounds' in order to determine how to restore the county to a full operational status."

Staff from Mullen Coughlin arrived in La Porte on Sunday night to assist. They will help prepare documentation to report the attack to the FBI and other appropriate law enforcement agencies.

Kora and Friedman praised the efforts of the IT Department.

Kora said: "I commend our IT Director Darlene Hale and her team for shutting down our systems Saturday afternoon as soon as the malware virus was detected. Unfortunately, at least half our servers have been infected and it will take some time to fully restore service. I ask for patience from the public as we seek to become fully operational again."

Friedman echoed that sentiment, saying: "Darlene Hale and her team have been working 15 hour days since this virus hit to try to restore portions of our system that can be restored. We ask for patience from all concerned."

So... that was the initial incident reporting...

Then this past Sunday, Bleeping Computer reported:

A forensic investigation firm and the FBI were involved, but attempts to recover the data encrypted by the malware without paying the ransom were fruitless.

The cybercriminals got about $130,000 in Bitcoin from this attack, with $100,000 being covered by insurance. The impact may not be immediate but it does create some ripples in the long run.

The decision to pay the cybercriminals came after seeing that the decryption keys from the FBI could not restore the encrypted files.

And according to a local report from WSBT, the county had backup servers… but the malware infected them.

We now know that insurance companies are bearing the brunt of the payouts for these attacks. So I'll bet that we're not far from the time when the conditions of continued insurance are regular training and reviews, periodic security audits and more reliable backup solutions. In other words: "We'll insure your municipality, but unless you want the insurance premiums to be sky-high, you need to get much more proactive about protecting yourself from these threats. And when you come calling for a payout, the first thing we do will be to audit to figure out why none of the multiple safeguards you promised to put in place -- and maintain -- were effective in this instance."

Meanwhile...

US mayors group adopts resolution not to pay any more ransoms to hackers
US mayors vow not to give in to more extortion demands following ransomware attacks.

The 2019 Adopted Resolutions of the 87th Annual Meeting of The United States Conference of Mayors, of the committee for Criminal and Social Justice, included the resolution to "oppose payment to ransomware attack perpetrators"

http://legacy.usmayors.org/resolutions/87th_Conference/proposedcommittee-preview.asp?committee=Criminal%20and%20Social%20Justice

"Opposing Payment To Ransomware Attack Perpetrators"

1. WHEREAS, targeted ransomware attacks on local US government entities are on the rise; and
2. WHEREAS, at least 170 county, city, or state government systems have experienced a ransomware attack since 2013; and
3. WHEREAS, 22 of those attacks have occurred in 2019 alone, including the cities of Baltimore and Albany and the counties of Fisher, Texas and Genesee, Michigan; and
4. WHEREAS, ransomware attacks can cost localities millions of dollars and lead to months of work to repair disrupted technology systems and files; and
5. WHEREAS, paying ransomware attackers encourages continued attacks on other government systems, as perpetrators financially benefit; and
6. WHEREAS, the United States Conference of Mayors has a vested interest in de-incentivizing these attacks to prevent further harm,
7. NOW, THEREFORE, BE IT RESOLVED, that the United States Conference of Mayors stands united against paying ransoms in the event of an IT security breach.

Uh huh. I was scanning through the much longer document, of which that is a small fraction, and I finally found the "Opposing Payment To Ransomware Attack Perpetrators" section. I suppose I thought that it was going to have somewhat more teeth. So... I it's nice that the nation's mayors "stand united against paying ransoms in the event of an IT security breach"... But so what?

I dug down into this a bit more and confirmed through other reporting that the resolution that was adopted last week doesn't have any legal binding, but that it can be used as an official position to justify administrative actions to both federal authorities and taxpayers. And, so yeah, that's all well and good. But when a city's systems have all been taken offline and non-ransom-payment remediation, if it's even possible, will apparently cost many more millions of dollars that post-payment remediation, and especially when a city has been paying insurance against this event, it seems to me that it all rather pragmatically comes down to money. Clearly, no one, other than the attackers, wants the ransom to be paid. But as we've covered recently when looking at the numbers, it's clear that those municipalities that choose to stand their ground and not pay probably ended up paying a great deal more in the end.

"Coveware" Ransomware Remediation
https://www.coveware.com/

- We are the first responders to your ransomware recovery.
- Coveware aggregates global ransomware data to minimize your ransomware related costs and downtime.
- Let our IT security professionals manage your ransomware incident response.

How do we help restore your encrypted data?

1 - Explore free remediation options
- Identify Ransomware Type
- Find free Decryptor tools
- Free initial risk assessment
- Identify Threat Actor Group

2 - Threat Actor Negotiations
- Secure & safe negotiations
- Complete & transparent communications
- Determine Risks & Outcomes

3 - Ransom Settlement
- 100% Transparency
- Reimbursed Costs
- Transparent Documentation
- Compliance Checks

4 - Restore data & end downtime.
- Professional IT support
- Insurance documentation
- Post-incident follow up, support

"Minimize your ransomware downtime. Let us manage your ransomware recovery."

"The Impossible Puzzle of Cybersecurity"
Results of an independent survey of 3,100 IT managers commissioned by Sophos
https://secure2.sophos.com/en-us/medialibrary/Gated-Assets/white-papers/sophos-impossible-puzzle-of-cybersecurity-wp.pdf

U.K.-based research house Vanson Bourne interviewed 3,100 IT decision makers between December 2018 and January 2019. To provide a representative size split within each country, respondents were split equally between 100-1,000 user organizations and 1,001-5,000 user organizations.

Respondents who had been victims of a cyberattack in the last year were asked how the most significant cyberattack got into their environment. The results revealed that, where respondents know how the attack got in, email is the most common attack vector, used in 33% of attacks. Given the prevalence of phishing, this is no surprise. The web is also a major vector, used in three in 10 attacks. Together, email and web account for nearly two-thirds of attacks entering organizations. IT managers can't just focus on email and web, however. 23% of attacks got in via a software vulnerability, and 14% via a USB stick or external device. Furthermore, 20% of IT managers didn't know how the most significant attack got in – if you don't know which security door has been left open it's hard to shut it.

#2 Cyberattacks are multi-stage, coordinated, and blended. Respondents whose organizations had been victims of a cyberattack revealed that they had suffered a wide range of attacks over the last year.

What type of cyberattack(s) has your organization been hit with
In The last year? Base: respondents from organizations that have
Fallen Victim to one or more cyberattack(s) in the last year (2109)

These numbers clearly add up to more than 100%, indicating that multi-stage attacks are now the norm. For example, a phishing email could install malicious code that takes advantage of a software exploit to install ransomware. The high numbers involved also confirms the scale of the challenge facing IT teams.

Phishing: the most prevalent cyberattack
Of the 2,109 organizations hit by a cyberattack in 2018, over half (53%) were victims of phishing. Indeed, phishing was also the most prevalent attack in all countries surveyed with the exception of Colombia, where it was the second most common threat. Across the full 3,100 respondents, over one-third (36%) fell victim to phishing emails.

Software exploits: varied impact around the globe
Of the organizations hit by a cyberattack, over a third (35%) suffered from an exploit taking advantage of a vulnerability in software they were using. There are significant regional variations in propensity to be affected by exploits. In Mexico, over half of the organizations that fell victim to a cyberattack experienced a software exploit (51%). This is more than double the number affected in Brazil (22%), South Africa, and Japan (both 23%)

#3 Technology, talent, and time are in short supply
As we've seen, organizations face a wide range of attacks and need to secure multiple threat vectors. The survey revealed that, on average, IT teams spend 26% of their time managing cybersecurity. For the majority of respondents this is not the right ratio.

Indian organizations spend the most time (32%) and Japanese teams the least (19%). Organizations that had been hit by a cyberattack spend a little more time on IT security (28%) than those that hadn't experienced an attack (23%).

Given the variety and complexity of threats, it's not surprising that 86% of respondents say they need greater cybersecurity skills in their organization. Those organizations that had experienced an attack have greater need for cybersecurity expertise than those that hadn't (89% vs. 79%). This could be because they have more security issues that need fixing, or the result of heightened awareness of the complexity of today's attacks.

However, bringing in the expertise to fill these gaps is a major challenge. Eight in 10 organizations say they struggle to recruit in the right skills. When it comes to recruitment, India faces the greatest challenge (89%) and Germany the least – but still, two in three German IT managers say they struggle to bring in the right skills.

At the same time, cybersecurity budgets are not sufficient with two in three (66%) respondents saying that their budget for people and technology is too low. This rises slightly to 70% in those organizations that were hit by a cyberattack in 2018.

The Fines are beginning to happen.

Mistakes are starting to cost more than just reputation... And I'm of two minds about fines: We **really** want major organizations to act responsibly with the personal and abusable data that they collect about us through the normal course of their justifiable business operations. But we also want -- and need -- them to self-report when, despite their best efforts, they fail to live up to their, and our, hopes. Given that responsible self-reporting is inherently voluntary, unless a breach is discovered externally, which is much less common than internal discovery, levying burdensome and abusive fines may not actually improve end-user security and privacy. So...

Marriott faces $123 million GDPR fine in the UK for last year's data breach
The UK's Information Commissioner's Office (ICO) has announced that it intends to impose a fine of £99,200,396 ($123,705,870) on the Marriott hotel chain over last year's data breach.

As we know and reported at the time, last November 2018, Marriott self-reported that hackers had access to the Starwood guest reservation database since 2014. Starwood being a chain that Marriott had acquired. Marriott initially reported that hackers stole the details of roughly 500 million hotel guests which they subsequently reduced to 383 million after a more throrough investigation. But, still... 383 million visitors. Ouch.

According to a post mortem of the hack, hackers stole:

- 383 million guest records
- 18.5 million encrypted passport numbers
- 5.25 million unencrypted passport numbers
- 9.1 million encrypted payment card numbers
- 385,000 card numbers that were still valid at the time of the breach

Class-action lawsuits began piling up within hours of Marriott's announced its security breach.

And I suppose not surprisingly, now the UK's Information Commissioner's Office which is in charge of such things, has stated that Marriott's security practices were in violation of the EU's GDPR -- the General Data Protection Regulation (GDPR).

https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2019/07/statement-intention-to-fine-marriott-international-inc-more-than-99-million-under-gdpr-for-data-breach/

Statement:

Intention to fine Marriott International, Inc more than £99 million under GDPR for data breach

Statement in response to Marriott International, Inc's filing with the US Securities and Exchange Commission that the Information Commissioner's Office (ICO) intends to fine it for breaches of data protection law.

Following an extensive investigation the ICO has issued a notice of its intention to fine Marriott International £99,200,396 for infringements of the General Data Protection Regulation (GDPR).

The proposed fine relates to a cyber incident which was notified to the ICO by Marriott in November 2018. A variety of personal data contained in approximately 339 million guest records globally were exposed by the incident, of which around 30 million related to residents of 31 countries in the European Economic Area (EEA). Seven million related to UK residents.

It is believed the vulnerability began when the systems of the Starwood hotels group were compromised in 2014. Marriott subsequently acquired Starwood in 2016, but the exposure of customer information was not discovered until 2018. The ICO's investigation found that Marriott failed to undertake sufficient due diligence when it bought Starwood and should also have done more to secure its systems.

Information Commissioner Elizabeth Denham said: "The GDPR makes it clear that organisations must be accountable for the personal data they hold. This can include carrying out proper due diligence when making a corporate acquisition, and putting in place proper accountability measures to assess not only what personal data has been acquired, but also how it is protected. Personal data has a real value so organisations have a legal duty to ensure its security, just like they would do with any other asset. If that doesn't happen, we will not hesitate to take strong action when necessary to protect the rights of the public."

Marriott has co-operated with the ICO investigation and has made improvements to its security arrangements since these events came to light. The company will now have an opportunity to make representations to the ICO as to the proposed findings and sanction.

The ICO has been investigating this case as lead supervisory authority on behalf of other EU Member State data protection authorities. It has also liaised with other regulators. Under the GDPR 'one stop shop' provisions the data protection authorities in the EU whose residents have been affected will also have the chance to comment on the ICO's findings.

The ICO will consider carefully the representations made by the company and the other concerned data protection authorities before it takes its final decision.

In a filing with the US Securities Exchange Commission last Tuesday, Marriott said it plans to appeal the ICO's fine, when formally filed. Marriott International's President and CEO, Arne Sorenson said: "We are disappointed with this notice of intent from the ICO, which we will contest. We deeply regret this incident happened. We take the privacy and security of guest information very seriously and continue to work hard to meet the standard of excellence that our guests expect from Marriott."

And Sorenson noted that Marriott had retired the compromised Starwood guest reservation system earlier this year.

This is the ICO's second announcement of plans to fine a large organization for GDPR violations in as many days. Last Monday, the ICO announced plans to fine British Airways £183 million ($230 million) after the company failed to protect its website, which was infected with a web-based card skimmer that collected payment details for British Airways customers between April and June 2018. So... it will be interesting to see how this turns out, and what/whether the GDPR-enabled aggressive fining activities help or hinder our security and privacy in the long run.

And speaking of fines, though along different lines, since in Facebook's case it was policy rather than a mistake…

Remember a few months ago when we talked about Mark Zuckerberg addressing his shareholders and stating that they had "set aside" (I think those were his words) some billions of dollars ("billions" with a "B") for an expected Federal Trade Commission fine in settlement of the infamous Cambridge Analytica-tied privacy violations? Well, the Wall Street Journal reported that FTC commissioners have voted to approach a $5 Billion settlement with Facebook.

Mozilla out of the UK's ISPA Doghouse

Paul Ducklin, a writer for Sophos "Naked Security", followed-up his earlier column about the nutty ISPA nomination of Mozilla as "Internet Villian of the Year" with a column titled: "Mozilla aren't villains after all" In his piece, Paul nicely summarizes why unprotected DNS over UDP is a problem in the first place:

"If I unlawfully sniff out your DNS traffic so I know where you went, I'm violating your privacy. Merely by knowing where you surfed, without getting any details of what you actually surfed, I can infer an awful lot about you. I can probably piece together your daily routine, both at work and at home; figure out your likes and fears; learn which companies you do business with; guess which bank you use, the shops you frequent, the clubs you belong to, the hobbies you enjoy, the medical surgery you're registered with, the sports teams you support, and much

more."

And as we all know, there are many other means for blocking access to unwanted sites. Since the DNS to IP mapping sometimes changes, an ISP's content blocking device could periodically make the same DNS queries their customers make, retrieve the DNS lookup IP, and dynamically add that to a IP-filter blocking list. Or some concerned organization could perform the lookups and communicate IP address additions and removals to concerned ISPs. Or ISPs could subscribe to a published "block list" in the same way as SPAM has been thwarted since 1997 with RBLs -- Realtime Black Lists -- of IPs of known spammers. There are a great many ways to solve this problem that are just as robust as filtering on DNS. Certainly those being filtered know that by changing their domain names they can sidestep the filtering until it catches up.

So, yeah, enhancing the privacy of ALL web browsing users at the expense of asking ISPs to change the details of the way they selectively block access so some domains (which haven't yet changed their names to avoid the blocking) makes a great deal of sense.

And speaking of Mozilla...

Recall that we have previously covered the shady organization "DarkMatter" which was petitioning Mozilla to include its CA root certificate in Firefox's trusted certs store. At the time, cyber-security experts and privacy advocates were strongly cautioning Mozilla against doing so stating that DarkMatter would abuse this position to help its surveillance operations. Some of these operations have been previously detailed in reports from Reuters, the New York Times, The Intercept, and other sources which detail alleged DarkMatter-orchestrated hacking operations against human rights activists, journalists, and foreign governments, which DarkMatter carried out at the behest of the United Arab Emirate' government.

And besides, if you pick the name "DarkMatter"?? Really?

In a last ditch effort to have its certificates trusted inside Firefox, DarkMatter most recently attempted to create a spin-off certificate authority business called "DigitalTrust". But -- whoopsie! -- DarkMatter and DigitalTrust were both run by the same CEO. These guys seem kinda clueless.

So... taking everything into consideration, giving plenty of time for consideration and contemplation, and really not want to deny anyone who SHOULD have this privilege out of hand, Mozilla announced its decision last week in a Google Groups discussion.

Wayne Thayer, Certificate Authority Program Manager at Mozilla, said: "Our foremost responsibility is to protect individuals who rely on Mozilla products. I believe this framing strongly supports a decision to revoke trust in DarkMatter's intermediate certificates. While there are solid arguments on both sides of this decision, it is reasonable to conclude that continuing to place trust in DarkMatter is a significant risk to our users. I will be opening a bug requesting the distrust of DarkMatter's subordinate CAs [...]. I will also recommend denial of the pending inclusion request, and any new requests from DigitalTrust."

The distrust of the subordinate CAs that Wayne was referring to was something we also discussed before. DarkMatter had been issuing certificates with an intermediate CA certificate signed by QuoVadis to obtain its trust. So that will be killed as well. Once Mozilla removes the QuoVadis intermediary certificates from Firefox in a future update, all websites that use TLS certificates acquired from DarkMatter will show full-page HTTPS errors in Firefox, warning and blocking users from accessing their content.

However, I'm wondering what this means for Windows certs? Recall that in order to prevent problems with 3rd party A/V, Mozilla will be importing the Windows CA root and trusting cart signed by those root certs. So if you are on Win8 or Win10 with a recent Firefox (since 66) and you have a non-Windows Defender A/V registered with the system, Firefox =may= be turning on the option to trust the Windows root store. If Windows is trusting DarkMatter, and/or their QuoVadic intermediate cert, then your Firefox would be, too.

about:config / security.enterprise_roots.enabled == FALSE

The next release of Firefox (69) adds a tracker blocking report...
The next release (#69) of Firefox will add a cool "about:protections" graphical display showing how many and of what type trackers Firefox has auto-blocked during the previous week.
https://mozilla.invisionapp.com/share/38S2U6FPFKD#/screens/360923276

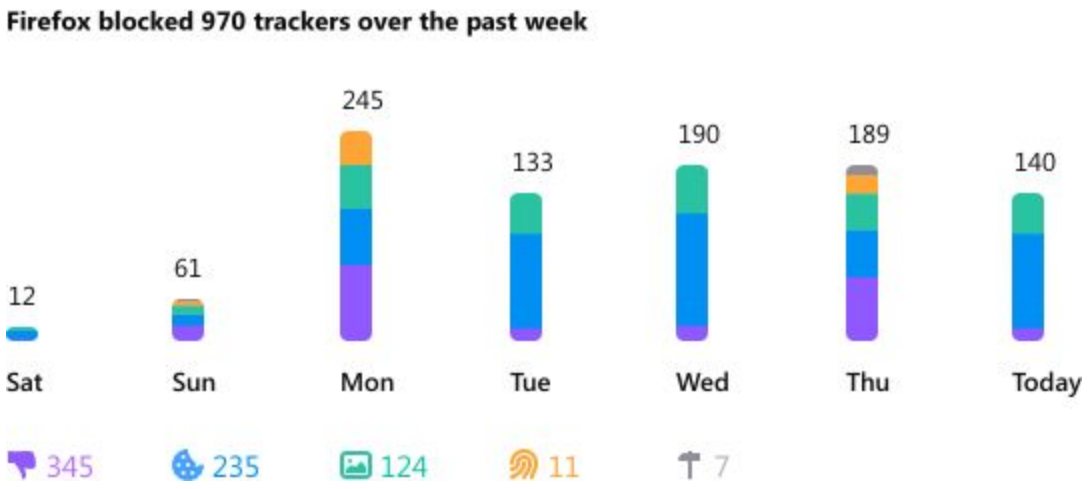

We didn't talk about this when it happened since we had so much news that week. But Mozilla has apparently clearly decided to differentiate itself from the Chromium-based browsers by focusing upon privacy through anti-tracking. They released the full version of its Enhanced Tracking Protection (ETP) system in Firefox 67.0.1 last month. It added default blocking for cross-site trackers, which are, as we know, small bits of JavaScript embedded in websites by advertising networks. Those bits of code send back our location to monitor what we're doing across the web for the purpose of generating profile of us.

At the same time, Mozilla released an updated version of its Facebook Container which stops Facebook from tracking people in a similar way. So those share and like buttons which appear

almost ubiquitously across the web to report back to Facebook even if they are never clicked are now completely blocked by the updated container, along with all other connections to Facebook's servers. And in May, Firefox also began blocking Cryptomining and browser fingerprinting.

Chrome is a great browser, and it now has the majority of the Internet. But we know how Google makes its money. I love their search engine and this Show Notes document was created using their very slick online tools. But I am more closely aligned with Firefox's philosophy, I love having tiny tabs in the sidebar, and Firefox works perfectly for me.

And now, Leo, in this week's installment of... "Wrestling a simple idea to the ground"
We have the paper which will be delivered this Thursday during the IEEE 43rd Annual Computer Software and Applications Conference (COMPSAC). It's titled: "CTRL-ALT-LED: Leaking Data from Air-Gapped Computers Via Keyboard LEDs"

My first thought upon hearing about this was that rather than being named "CTRL-ALT-LED" it should be called "CTRL-ALT-DUH"... because it's utterly obvious to all of us that if software can blink a keyboard's LEDs, and software can, and malicious software can be pre-installed into a computer -- which is a prerequisite for this technique -- then, if there is also some way to arrange to have something watching the keyboard's LEDs, which is, of course, required -- and presumably when no one is around, since they would notice that something screwy was going on if their keyboard's LEDs suddenly went berserk -- that, indeed, it would definitely be possible to exfiltrate data from that computer.

We've covered their work before. They appear to really like blinking lights. They have previously described how hard drive LEDs could be used to exfiltrate data, and also how router hardware, if the lights actually blinked with the data -- which is typically not the way those lights blink -- could also be used to exfiltrate data. So, like I said, these guys really like their blinking light.

To their credit, when they tackle something that's kinda obvious to us all, they do, at least, as I noted above, really wrestle the subject to the ground. Because they wanted to answer the question, not whether it was possible -- which, given all of the prerequisites noted, is obvious to us all -- but exactly how quickly said data could be extracted given every available trick in the book. So the Abstract of their paper reads:

Abstract:
Using the keyboard LEDs to send data optically was proposed in 2002 by Loughry and Umphress [1] (Appendix A). In this paper we extensively explore this threat in the context of a modern cyber-attack with current hardware and optical equipment. In this type of attack, an advanced persistent threat (APT) uses the keyboard LEDs (Caps-Lock, Num-Lock and Scroll-Lock) to encode information and exfiltrate data from airgapped computers optically. Notably, this exfiltration channel is not monitored by existing data leakage prevention (DLP) systems. We examine this attack and its boundaries for today's keyboards with USB controllers and sensitive optical sensors. We also introduce smartphone and smartwatch cameras as components of malicious insider and 'evil maid' attacks. We provide the necessary scientific background on optical communication and the characteristics of modern USB keyboards at the hardware and software level, and present a transmission protocol and modulation schemes. We implement the exfiltration malware, discuss its design and implementation issues, and evaluate it with different

types of keyboards. We also test various receivers, including light sensors, remote cameras, 'extreme' cameras, security cameras, and smartphone cameras. Our experiment shows that data can be leaked from air-gapped computers via the keyboard LEDs at a maximum bit rate of 3000 bit/sec per LED given a light sensor as a receiver, and more than 120 bit/sec if smartphones are used. The attack doesn't require any modification of the keyboard at hardware or firmware levels.

https://ieeexplore.ieee.org/document/8754078

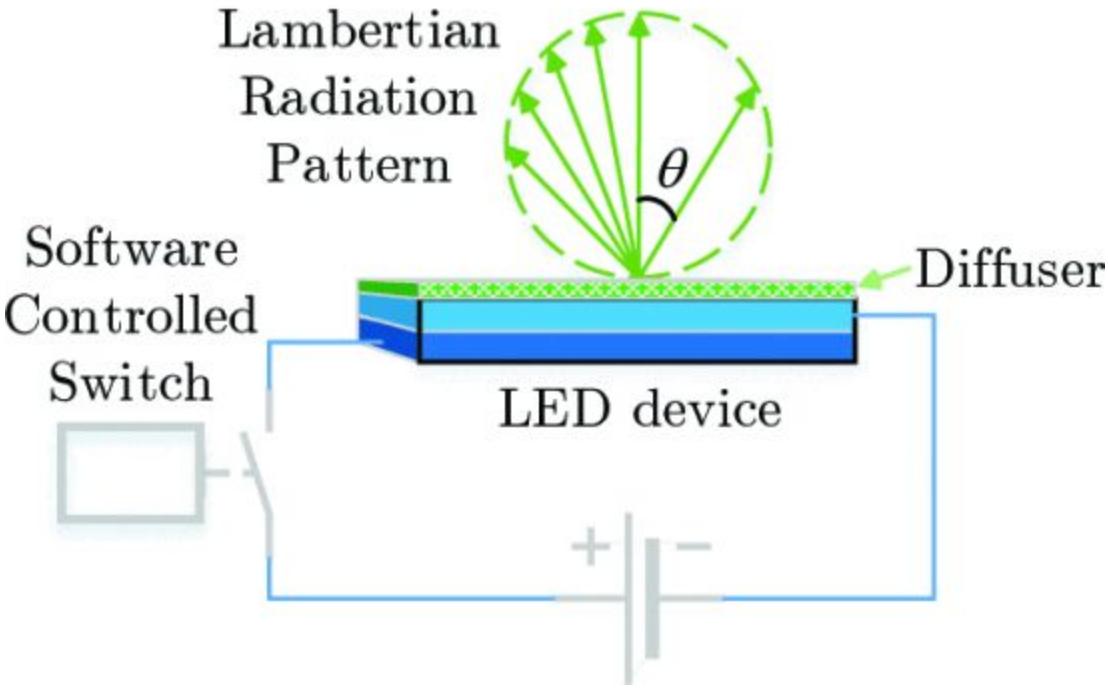

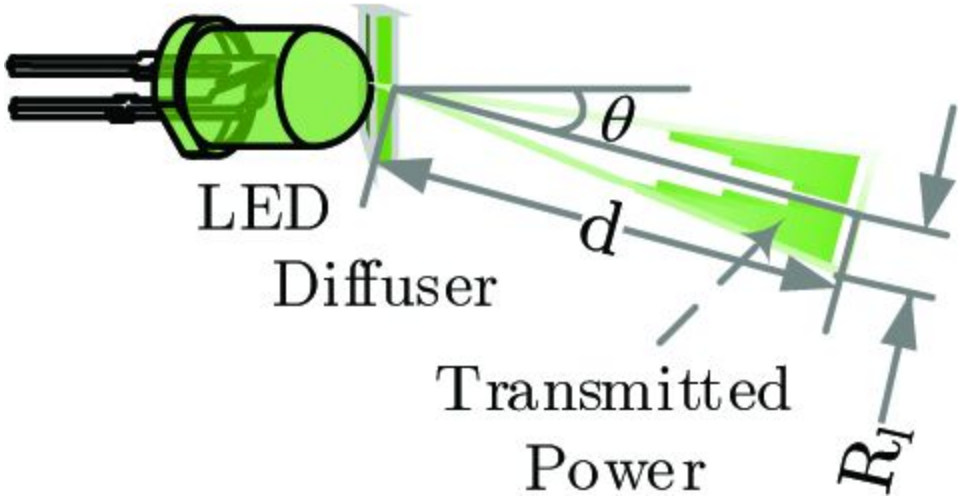

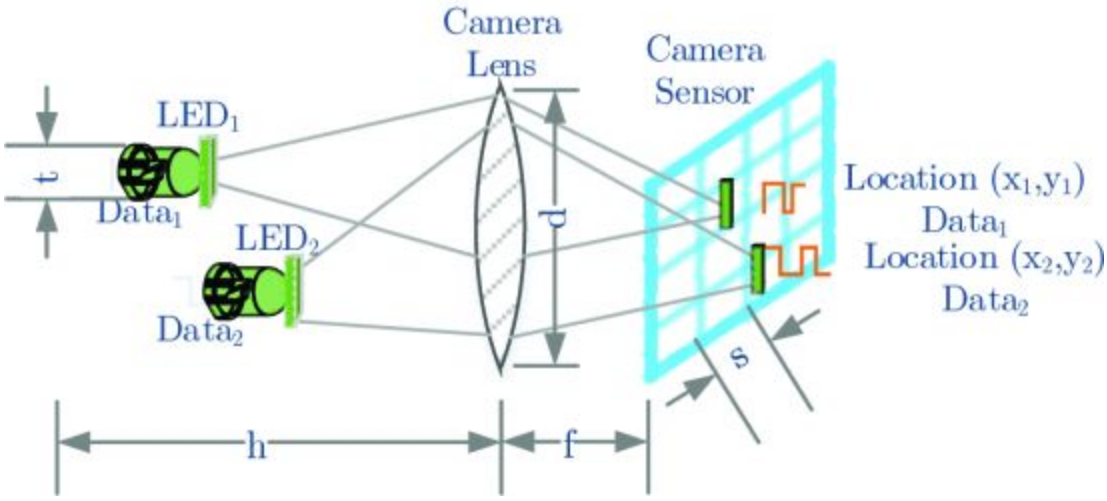

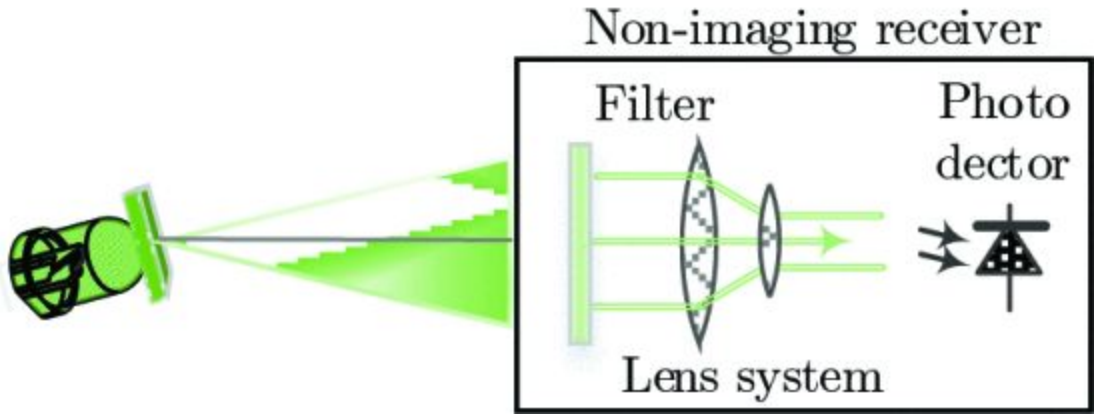

Keyboard	Modulation	Bit-rate	BER in %
Dell	OOK	1666 bit/sec	3%
Dell	Multiple LEDs	3411 bit/sec	2.40%
Lenovo	OOK	2230 bit/sec	2.95%
Lenovo	Multiple LEDs	4640 bit/sec	6.70%
Logitech	OOK	2170 bit/sec	3.50%
Logitech	Multiple LEDs	4296 bit/sec	1.20%
Silverline	OOK	2697 bit/sec	8%
Silverline	Multiple LEDs	5155 bit/sec	3.10%

Errata

Elaine: Oh, almost forgot, Year 15 begins August 20.

Anonymous Sender / Location: Somewhere
Subject: SecurityNow dnscrypt is compatible with DNSSEC
Date: 15 Jul 2019 17:57:07
:
I remember you or Leo saying that dnscrypt is incompatible with DNSSEC. I use dnscrypt and DNSSEC seems to be working according to https://dnssec.vs.uni-due.de/

Miscellany

"Stranger Things 3" was a disappointment.
Blatant product placement. In one shot of a supermarket isle, one side was crisply in focus showing all cereal products of one vendor while the opposing side of the isle was unreadably blurred out. And the venue being "The Mail" was such a blatant commercialization.

Closing The Loop

"Chuck" posted to GRC's Security Now! newsgroup:
I enjoyed show #722 yesterday.

Last night I checked the Firefox setting to Enable DNS over HTTPS (DOH). I chose a VPN location in Europe and started browsing.

There was a dramatic improvement in the speed with which websites started loading on the screen. I mean really fast.

Tonight I'm going to do some unchecking and rechecking to confirm that DOH is responsible.

Is DOH improving the efficency of a VPN connection?

Chuck followed-up exactly 24 hours later, with:
I disabled DNS Over HTTPS (DOH) in Firefox last night. Web page loading performance slowed considerably.

Naturally I turned it back on and the joy of quickly loading web pages returned!

Encrypting DNS

So... we have DNSSec, DNSCrypt, DNS over HTTPS (DoH) and DNS over TLS (DoT)

DNSSec...
...provides cryptographically signed DNS records which allows a DNSsec-aware OS to verify that the DNS response received, as it has been cached and forwarded from its authoritative DNS serving originator, has not been tampered with or altered in any way. Since the DNS reply is signed with a private key which no forger can have, this essentially means that we are assured that the received DNS reply is authentic.

That's all good. But what DNSSec does NOT do is encrypt. It was never intended to provide privacy, only authenticity. So the records are signed and cannot be tampered with, but anyone watching the traffic will see the DNS client's queries and their replies just as if DNSSec was not in use.

Before I go on I'll note that ALL of the three full encryption options are 100% compatible with DNSSec. The earliest versions of DNSCrypt were NOT compatible with DNSSec, but that has not been true for some time. So DNSSec CAN be used by any and all of the three following DNS encrypting protocols.

We first discussed DNSCrypt back in the context of OpenDNS (now owned by Cisco). DNSCrypt used the same fast, lean and secure crypto that I chose for SQRL -- Dan Bernstein's Elliptic Curve 25519. It successfully provides encryption for privacy, but it is not nearly as attack and hack resistant as we would wish for a contemporary protocol since it doesn't use any of the existing public certificate infrastructure. The server's public key is published over DNS and is implicitly trusted, though it can be verified with DNSSec. So this meant that DNSCrypt was simple and lightweight, and that it could ride atop either UDP or TCP. Unlike any connection-oriented protocol (DoH or DoT) it requires much lighter server resources, it doesn't require a TLS stack and the security troubles that can bring. But it doesn't have an RFC and was never taken up by the IETF for Internet standardization.

That's why I've been using the past tense when referring to DNSCrypt. It was a pioneer in DNS encryption, but the properties of DNS over TLS or HTTP, while lacking some of DNSCrypt's laudable features, have ended up winning the day.

But what's nice is that there are DNSCrypt clients which have been extended to support DoH

The reference appears to be "DNSCrypt-Proxy" written by Frank Denis (@jedisct1) in GoLang, supports both DNSCrypt and DoH and provides DNS client services for Linux, BSD, Windows, macOS, Android and more.
https://github.com/jedisct1/dnscrypt-proxy

And "Simple DNSCrypt" is a very nice and friendly configuration frontend for Windows written in C#. https://simplednscrypt.org/

DNS over TLS (DoT)

DNS over TLS is obviously a protocol for encrypting and wrapping DNS queries and their replies in TLS. This offers both privacy via TLS' encryption and authentication via TLS' support for the entire public key infrastructure. So this prevents eavesdropping and any manipulation of DNS data via man-in-the-middle attacks which, as we know, simply DNS over UDP is extremely prone to.

Cloudflare, IBM's Quad9, Google, Quadrant Information Security and CleanBrowsing are providing public DNS resolver services via DNS over TLS. Back in April of 2018, Google announced that Android Pie will include support for DNS over TLS. DNSDist, from PowerDNS also announced support for DNS over TLS in its latest version. Users of the older BIND DNS server users can also provide DNS over TLS by proxying it through stunnel. The newer Unbound DNS server has supported DNS over TLS natively since early last year.

So DNS over TLS is a nice option, especially is your client platform -- like Android Pie -- can support it natively. Just turn it on!
https://developers.cloudflare.com/1.1.1.1/setting-up-1.1.1.1/android/

1. Go to Settings ? Network & internet ? Advanced ? Private DNS.
2. Select the Private DNS provider hostname option.
3. Enter one.one.one.one or 1dot1dot1dot1.cloudflare-dns.com and hit Save.
4. Visit 1.1.1.1/help to verify DNS over TLS is enabled.

Also, last week I misspoke about the pfSense firewall support. pfSense supports DNS over TLS, not HTTPS.

Despite all that, DNS over TLS feels as though it has been eclipsed by DNS over HTTPS which appears to have a bit more momentum behind it…

DNS over HTTPS

It is a proposed IETF standard, as I mentioned last week, specified under RFC 8484. It uses HTTP/2 and HTTPS, and supports the on-the-wire format of DNS response data, as returned in existing UDP responses, so it's extremely simple to bring up on a web server. It defines a new HTTPS payload with the MIME type application/dns-message. When HTTP/2 is used, the server may also use HTTP/2 server push to send values that it anticipates the client may find useful in advance.