Security Now! #871 - 05-17-22

The New EU Surveillance State

This week on Security Now!

This week we look back at what no one wanted: an eventful patch Tuesday. Apple has pushed a set of updates to close an actively exploited 0-day. Google announced the creation of their Open Source Maintenance Crew. A ransomware gang wants to overthrow a government. Google's Play Store faces an endlessly daunting task. The predicted disaster for F5's BIG-IP systems arrived. A piece of errata and some closing the loop feedback from our terrific listeners. Then we're going to look at just how far afield the European Union has wandered with their forthcoming breathtaking surveillance legislation.

A perfect snapshot to characterize the current state of the security of the open source software supply chain:

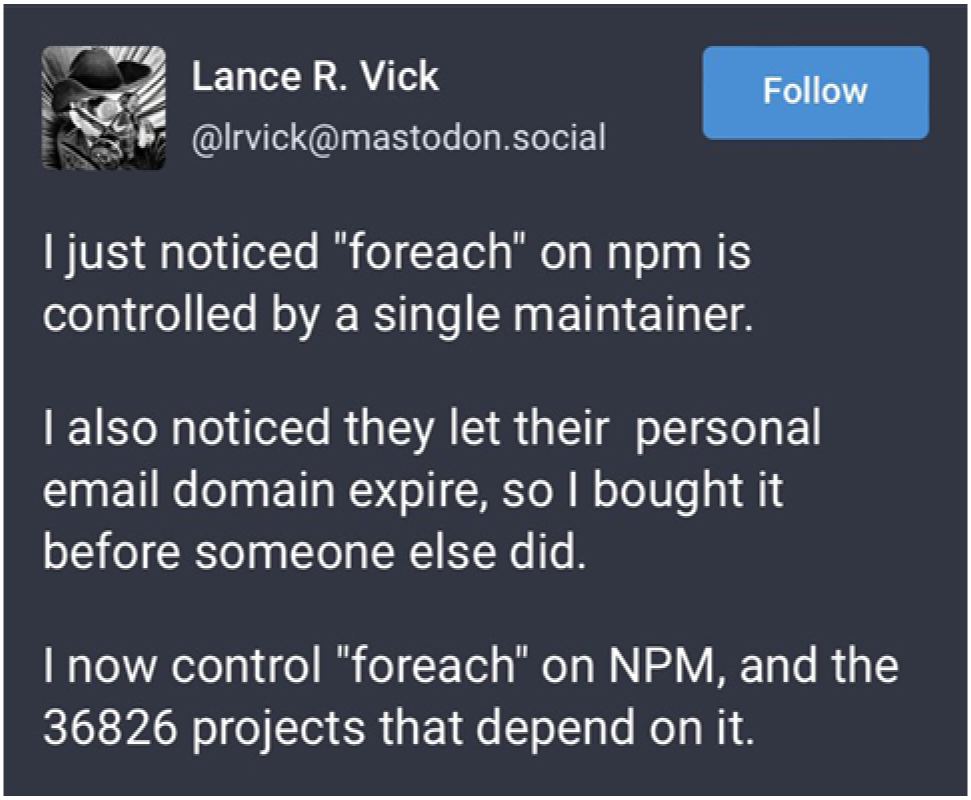

Lance R. Vick
@lrvick@mastodon.social

I just noticed "foreach" on npm is controlled by a single maintainer.

I also noticed they let their personal email domain expire, so I bought it before someone else did.

I now control "foreach" on NPM, and the 36826 projects that depend on it.

The repository change log shows that the project's eMail address domain was changed 6 days ago. We'll be discussing the OpenSSF — The Open Source Security Foundation again today ... And not a moment too soon!

Security News

An “eventful” Patch Tuesday

I've observed before that what one looks for in a Patch Tuesday is a seamless and uneventful experience. Either you check for updates and decide to install them, or you receive a notice that updates have arrived and are already installed and are just waiting for you to step away from your computer. Again, uneventful. What you don't want is to see a headline such as BleepingComputer ran yesterday: *“CISA warns not to install May Windows updates on domain controllers.”*

In fact, CISA so much doesn't want May's updates installed, that they went so far as to temporarily REMOVE the listing of one of their MUST PATCH security flaws from its catalog of known exploited vulnerabilities because they really can't have it listed there while they're also warning all users of Active Directory not to install those updates to fix it.

https://www.cisa.gov/uscert/ncas/current-activity/2022/05/13/cisa-temporarily-removes-cve-2022-26925-known-exploited

The headline on CISA's published notice reads: *“CISA Temporarily Removes CVE-2022-26925 from Known Exploited Vulnerability Catalog”*

“CISA is temporarily removing CVE-2022-26925 from its Known Exploited Vulnerability Catalog due to a risk of authentication failures when the May 10, 2022 Microsoft rollup update is applied to domain controllers. After installing May 10, 2022 rollup update on domain controllers, organizations might experience authentication failures on the server or client for services, such as Network Policy Server (NPS), Routing and Remote access Service (RRAS), Radius, Extensible Authentication Protocol (EAP), and Protected Extensible Authentication Protocol (PEAP). Microsoft notified CISA of this issue, which is related to how the mapping of certificates to machine accounts is being handled by the domain controller.”

It was once possible to individually install security patches so that a troublesome patch could be manually avoided. But now everything is rolled up into a take-them-all or none solution. And I don't blame Microsoft for that. My mind was always boggled that Micrsosoft was even able to consider offering a la carte patching of such an incredibly complicated codebase as Windows has become. The complexity of offering that option was astonishing. And, as we know, it often didn't quite work as hoped. So now it's all or nothing. And in the case of Windows domain controllers, for the time being, **“nothing”** is what you want.

As enterprise admins began installing the May updates last week, problems quickly started surfacing with admins sharing reports of some Active Directory policies failing with the error message: “Authentication failed due to a user credentials mismatch. Either the user name provided does not map to an existing account or the password was incorrect.”

Microsoft explained that the issue is only triggered after installing the updates on servers used as domain controllers. The updates will not negatively impact when deployed on client Windows devices and non-domain controller Windows Servers. And this is an example of an instance where we will eventually learn whether Microsoft's announced and forthcoming “AutoPatch” system is a good thing or more trouble than it's worth. Presumably, AutoPatch is somehow going to handle unforeseen problems like this. At the moment it's unclear where this omniscience is going to come from... Since it's apparently not coming from Microsoft.

The problem all this is trying to fix, in this case, is a flaw that's being actively exploited in the wild. It's a Windows LSA (Local Security Authority) spoofing 0-day which has been confirmed as a new PetitPotam Windows NT LAN Manager Relay attack vector. The PetitPotam problem was discovered and named last July by the French security researcher Gilles (je’-ill-a) Lionel. We talked about it at the time, which is probably why it sounds familiar. An NTLM Relay Attack allows bad guys to force devices, in this case domain controllers, to authenticate against malicious servers they control — essentially joining the malicious server to the domain. Once a device authenticates, the malicious server can impersonate the device and gain all of its privileges. This, in turn, gives attackers complete control over the domain.

As for the actual patch itself, we’re back to that disheartening story of Microsoft patching to stop a proof of concept from functioning, while leaving the underlying problem unresolved. Gilles has confirmed to the tech press that May's security update has finally fixed the specific problem, again, after he discovered a simple work-around for Microsoft’s first attempt to fix it last August. Yet even now, because the underlying problem remains, Gillles said that other EFS (encrypted file system) attack vectors still exist which will allow a slightly modified attack to continue to work. He said: “All functions of PetitPotam, as other vectors, still work except EfsOpenFileRaw.”

All of this is current as of yesterday. So assuming that Microsoft is eventually able to fix and reissue May’s security bundle in a way that doesn’t break Active Directory servers—and hopefully they’ll test it first to be sure—then they will either announce an out-of-schedule update or perhaps they’ll wait until June.

Patch Tuesday

We know what some Windows Domain Controllers **didn't** get last Tuesday. But what **did** most of us get? We received fixes for three new 0-day vulnerabilities and patches for a total of 75 flaws in Microsoft's software, eight of which are rated “Critical.”

- 26, one more than one third of the 75 flaws were Remote Code Execution Vulnerabilities
- 21 were Elevation of Privilege Vulnerabilities
- 17 Information Disclosure Vulnerabilities
- 6 Denial of Service Vulnerabilities
- 4 Security Feature Bypass Vulnerabilities
- 1 Spoofing Vulnerability
- 0 Edge - Chromium Vulnerabilities

Remember that Microsoft classifies a flaw as a 0-day if it's been disclosed publicly whether or not it's known to be actively exploited. In this case, of the three 0-days, only that worrisome PetitPotam NTLM relay attack is known to be exploited.

Given that “Exploit Wednesday” now follows “Patch Tuesday”, the urgency to install updates in a timely manner has increased.

Apple patched a 0-day

Yesterday, Apple updated watchOS to v8.6, tvOS to v15.5 and macOS Big Sur to v11.6. So, Apple Watch Series 3 or later, Apple TV 4K 1st and 2nd gen, Apple TV HD and Macs running Big Sur. In each of those cases the updates fixed an out-of-bounds write which could be made to occur in the AppleAVD module, which is a kernel extension for handling audio and video decoding. This will come as no surprise to our long-time listeners who will have all learned to expect trouble to arise in complex media decoders which are inherently complex interpreters of encoded bitstreams. In this case, remote attackers could have — and were known to be — executing their arbitrary code with kernel privileges. Apple was as closed mouthed as usual about this, only saying that they had added improved bounds checking.

Google's "Open Source Maintenance Crew"

Recall that two weeks ago we first talked about the "OpenSSF" — the Open Source Security Foundation. At that time I enumerated the gratifyingly large number of participating and supporting companies, and the occasion was their announcement of the Package Analysis Project which, in just one month, had identified more than 200 malicious packages which were present in the Python and JavaScript repositories. And recall that I was a bit wary of getting too excited about this particular effort since it appeared that they were mostly just scanning for references to previously known malicious domains and IPs... which would all be trivial to change once it became clear to the bad guys that this was the way to avoid this particular detection.

But as for the OpenSSF effort overall, I'm very bullish about the prospect of this. It's what has been needed for so long. Open source software originated as a counter culture phenomenon back in the days when source code was not commonly shared for any purpose and the idea of doing so was bizarre. Back then, the idea of software being free represented a clear threat to the interests of commercial proprietary software vendors. In February of 2001, Microsoft's Jim Allchin publicly stated that "open source is an intellectual property destroyer. I can't imagine something that could be worse than this for the software business and the intellectual-property business." And early the following year, in January of 2002, one of Microsoft's chief strategists, Craig Mundie addressed New York University's School of Business, saying that releasing source code into the public domain is "unhealthy", causes security risks and "as history has shown, while this type of model may have a place, it isn't successful in building a mass market and making powerful, easy-to-use software broadly accessible to consumers."

Now, that was then, and no one is holding Microsoft responsible for anything said 20 years ago. The world is an entirely different place today. But it does remind us just how much things have changed in 20 years. And we know that change is slow. We also know that the Open Source model has produced tremendous wealth — both intellectual and economic — and it has become a crucial component of today's software technology landscape, which even Microsoft has now begun to embrace. Today, it is entirely possible to operate a major enterprise using nothing but open source software.

But its problems are many, too. The trouble is that volunteer effort is more interested in creating than in maintaining and securing. It's not that maintenance and security focus are absent, but as we have seen, so much maintenance and security focus is needed beyond just getting something to work that it's a big ask. And truly securing software, understanding the many ways in which code which works can still be made not to work, requires a different mindset and a very different type of very specific education and training.

So many major organizations are now benefiting from the work that has been done for them, that having them join a Foundation so that they have an organized platform for giving something back — especially when it's about improving the crucial security of the software they are now all using within their enterprises and on their network borders — is the right thing to do. And the OpenSSF is that foundation.

We're talking about this again today because last Thursday Google made a major announcement of specific new support for this effort:

https://blog.google/technology/safety-security/shared-success-in-building-a-safer-open-source-community/

They wrote:

> *Today we joined the Open Source Security Foundation (OpenSSF), Linux Foundation and industry leaders for a meeting to continue progressing the open source software security initiatives discussed during January's White House Summit on Open Source Security. During this meeting, Google announced the creation of its new* ***"Open Source Maintenance Crew"*** *— a dedicated staff of Google engineers who will work closely with upstream maintainers on improving the security of critical open source projects. In addition to this initiative, we contributed ideas and participated in discussions on improving the security and trustworthiness of open source software.*
>
> *Amid all this momentum and progress, it is important to take stock on how far we've come as a community over the past year and a half. In this post we will provide an update on some major milestones and projects that have launched and look towards the future and the work that still needs to be done.*
>
> *A little over a year ago we published* ***Know, Prevent, Fix,*** *which laid out a framework for how the software industry could address vulnerabilities in open source software. At the time, there was a growing interest in the topic and the hope was to generate momentum in the cause of advancing and improving software supply-chain security.*
>
> *The landscape has changed greatly since then:*
>
> - *Prominent attacks and vulnerabilities in critical open source libraries such as Log4j and Codecov made headline news, bringing a new level of awareness to the issue and unifying the industry to address the problem.*
>
> - *The US government formalized the push for higher security standards in the May 2021 Executive Order on Cybersecurity. The release of the Secure Software Development Framework, a set of guidelines for national security standards on software development, sparked an industry-wide discussion about how to implement them.*

> - *Last August, technology leaders including Google, Apple, IBM, Microsoft, and Amazon invested in improving cybersecurity — and Google alone pledged* ***$10 billion over the next five years*** *to strengthen cybersecurity, including* ***$100 million to support third-party foundations, like OpenSSF****, that manage open source security priorities and help fix vulnerabilities.*
>
> *In light of these changes, the* ***Know, Prevent, Fix*** *framework proved prescient: beyond just the increased discussion about open source security, we're witnessing real progress in the industry to act on those discussions. In particular, the OpenSSF has become a community town hall for driving security engineering efforts, discussions, and industry-wide collaboration.*

Google's post goes into greater details about their plans for participation. But I just wanted to follow-up on our introduction of the OpenSSF two weeks ago to note that this is looking like the organization that's going to succeed. Previous efforts were well-meaning but premature, and as history shows, visionaries are often too far ahead of the pack.

It feels like the open source movement is finally being recognized and is earning the respect it deserves. It may have taken something like the scare of the Log4j vulnerability, giving major corporations a bit of a wake up call, to realize just how dependent they had grown on open source solutions through the years. But either way, it appears that is finally happening now.

Conti suggests overthrowing the new Costa Rican government

As I've promised, I won't spend lots of our listeners' valuable time discussing boring details of endless ransomware attacks. But when a ransomware gang gets so big for their britches that they suggest that perhaps a government which is refusing to pay their ransom should be overthrown by its citizenry, I think that rises to a new level of interest.

I referred to this drama for the first time, mostly in passing, last week. Russia's Conti ransomware gang is behind the attacks on several Costa Rican government ministries. Over the weekend they doubled their ransom demand from $10 million to $20 million.

The Costa Rican operations which have been affected are:

- The Finance Ministry
- The Ministry of Science, Innovation, Technology, and Telecommunications
- The Labor and Social Security Ministry
- The Social Development and Family Allowances Fund
- The National Meteorological Institute
- The Costa Rican Social Security Fund
- The Interuniversity Headquarters

In two messages posted to Conti's leak site Saturday, the gang which has already leaked 97% of the 670 GB stolen during their attacks, claimed the U.S. government, was <quote> "sacrificing" Costa Rica and that the country's government should pay for the decryption keys to unlock their systems.

As I mentioned last week, Costa Rica's new government had taken office just last week and immediately declared a state of emergency after refusing to pay the initial $10 million ransom demand issued by Conti. Costa Rica has received assistance from officials in the U.S., Israel and other countries. And the context for me mentioning all this last week was the U.S. State Department's announcement of a $10 million bounty for information about anyone connected to Conti with an additional $5 million payable for information leading to an arrest and conviction.

Conti posted:

> *"Why not just buy a key? I do not know if there have been cases of entering an emergency situation in the country due to a cyber attack? In a week we will delete the decryption keys for Costa Rica."*

> *"I appeal to every resident of Costa Rica, go to your government and organize rallies so that they would pay us as soon as possible. If your current government cannot stabilize the situation? Maybe it's worth changing it?"*

Like I said, too big for their britches. In another message, the group called President Joe Biden a "terrorist" (probably as a result of the State Department's new bounty declaration) and said it was raising the ransom to $20 million. The group also implied that it would begin calling government officials to demand the ransom. (Yeah, like they've got a spare $20 million in their pockets.)

> *"Just pay before it's too late, your country was destroyed by 2 people, we are determined to overthrow the government by means of a cyberattack, we have already shown you all the strength and power, you have introduced an emergency."*

It's true that Costa Rica is limping along at the moment. The attack crippled the country's customs and taxes platforms alongside several other government agencies, even bringing down one Costa Rican town's energy supplier. The country's treasury department has been unable to operate any of its digital services since the attack began, making it nearly impossible for paperwork, signatures and stamps — required by law — to be processed. More than three weeks after the attack began, the country is still facing significant struggles, particularly because of the damage done to the Finance Ministry. Last week the country told residents that taxes need to be calculated by hand and paid in person at local banks, as opposed to the digital system the country has previously used.

Policing the Google Play Store

There's a probably-intractable problem with the model we currently have for freely downloadable mobile device apps created by individuals with no reputation. After all, everyone starts off having no reputation. Android handsets are available for a fraction of the price of Apple's devices and Android users typically cite the expansive freedom provided by the Android platform as their primary reason for preferring that much more open mobile environment. But those listening to this podcast realize that with that freedom comes significantly increased danger.

I think it’s clear that Google is doing the best they can to minimize this danger. But a continuous daily incoming torrential flood of apps is arriving at the Play Store, and there is just no way for Google to deeply research the behavior of each and every one of those apps. And to provide the useful and powerful freedom that Android users demand, apps must be given powerful enough access to the underlying hosting platform that a malicious app could be quite abusive. So Google is always stuck playing catch up, and in addition to their own efforts, relies upon the motivations and scrutiny that’s also offered by 3rd party security companies.

So, yesterday, Trend Micro posted their piece titled “Fake Mobile Apps Steal Facebook Credentials & Cryptocurrency-Related Keys.” In their article, Trend Micro explained that malware that's expressly designed and intended to steal the Facebook logon credentials of Android phone users continues to pop up on the Google Play Store. Such malware has become so commonplace that it's being called “Facestealer” malware. But it doesn't say that on the cover. It's hidden in apps that otherwise look harmless, compelling and, of course, free.

Trend Micro recently identified more than 200 Facestealer variants in the store, notified Google, and Google took them down. But how long will it be before they're replaced by 200 more? Some of the apps that were just taken down had been installed more than one hundred thousand times. The apps take the form of tools for editing, manipulating or sharing photos, but they can take other forms.

An example was “Daily Fitness OL” which appears to be a fitness app complete with exercises and video demonstrations. But it was designed to steal the Facebook credentials of its users. These so-called Facestealer apps were first identified in July of last year and have been linked to Russian servers by researchers with the mobile security company Pradeo. Attackers typically use the compromised Facebook accounts they acquire for various malicious purposes such as phishing scams, fake posts, and ad bots.

In the case of “Daily Fitness OL”, users are prompted to log in to Facebook through an embedded browser (what could possibly go wrong with that?), then a piece of JavaScript is injected into the loaded webpage which, of course, steals the credentials entered by the user. Easy peasy.

Trend Micro identified many other Facestealer apps with names like Enjoy Photo Editor, Panorama Camera, Photo Gaming Puzzle, Swarm Photo and Business Meta Manager. And in addition to these 200+ Facestealer apps, Trend Micro noted that they had found about 40 fake cryptocurrency mining apps that are designed to steal their user's cryptocurrency.

Last month Google reported that last year they had removed more than 1 million malicious apps from the Play Store. Think about that! 1 million malicious apps in 2021. An intractable problem.

And the trouble is, there is next to zero general awareness of this problem among the Android using population. There are presently more than 3 billion - with a 'B' - active Android devices being used worldwide. There's no question that the majority of Google Play Store apps are legitimate and well meaning. But when a malicious app is only removed after having been downloaded and installed more than one hundred thousand times, it's also clear that downloading Android apps carries a non-zero risk... and it's not clear that anything can be done about that.

The situation has grown more dire for F5 systems' BIG-IP boxes.
The day after we talked about this last week CISA added that recently disclosed F5 BIG-IP flaw to its Known Exploited Vulnerabilities Catalog following reports of active abuse in the wild.

The problem is CVE-2022-1388, bearing a well-deserved CVSS of 9.8, due to a critical bug in BIG-IP's iControl REST endpoint which provides an unauthenticated attacker with a method to execute arbitrary system commands. The firm Horizon3.AI wrote: *"An attacker can use this vulnerability to do just about anything they want to on the vulnerable server. This includes making configuration changes, stealing sensitive information and moving laterally within the target network."*

Although patches and mitigations for the flaw were announced by F5 on May 4th, the Wednesday before last, we know how well that tends to go. And, in fact the F5 boxes have been subjected to in-the-wild exploitation ever since F5's announcement with some attackers attempting to install a web shell that grants backdoor access to the targeted systems and others simply attempting to destroy the device's usability by executing a recursive "rm" — remove all files — starting from the device's root directory.

Rapid7 wrote: *"Due to the ease of exploiting this vulnerability, the public exploit code, and the fact that it provides root access, exploitation attempts are likely to increase"* but their security researcher Ron Bowes added that: *"Widespread exploitation is somewhat mitigated by the [relatively] small number of internet-facing F5 BIG-IP devices."*

The SANS Internet Storm Center (ISC) wrote on Twitter that *"Given that the web server runs as root, this should take care of any vulnerable server out there and destroy any vulnerable BIG-IP appliance."*

Pursuant to CISA's addition of this vulnerability to their catalog, all Federal Civilian Executive Branch agencies have been mandated to patch all systems against this issue by May 31st, two weeks from today. Of course, by that time, there will be nothing left standing to patch.

Errata

Since it's an interesting and important topic that's perfect for this podcast, I want to take a moment to talk about classical computing, quantum computing and symmetric vs asymmetric cryptography.

Back in 1994, an American mathematician by the name of Peter Shor conceived of an algorithm for quantum computers which would be able to determine the prime factors of integers. That algorithm worked and it bears the name "Shor's Algorithm."
https://en.wikipedia.org/wiki/Shor%27s_algorithm

Wikipedia explains that: *"The efficiency of Shor's algorithm is due to the efficiency of the quantum Fourier transform and modular exponentiation by repeated squarings. If a quantum computer with a sufficient number of qubits could operate without succumbing to quantum noise and other quantum-decoherence phenomena, then Shor's algorithm could be used to break public-key cryptography schemes, such as*

• The RSA scheme
• The Finite Field Diffie-Hellman key exchange
• The Elliptic Curve Diffie-Hellman key exchange"

In other words, I was incorrect to state last week that the use of elliptic curve crypto was "post-quantum" safe. It's generically any **asymmetric public-key crypto** that isn't safe. And I know better, so I wanted to correct the record. It's **symmetric crypto** that remains safe in a post-quantum crypto world.

Wikipedia explains: *"RSA is based on the assumption that factoring large integers is computationally intractable. As far as is known, this assumption is valid for classical (non-quantum) computers; no classical algorithm is known that can factor integers in polynomial time. However, Shor's algorithm shows that factoring integers is efficient on an ideal quantum computer, so it may be feasible to defeat RSA by constructing a large quantum computer. It was also a powerful motivator for the design and construction of quantum computers, and for the study of new quantum-computer algorithms. It has also facilitated research on new cryptosystems that are secure from quantum computers, collectively called post-quantum cryptography."*

The good news is, from a practical standpoint it still looks like we're well away from the quantum crypto apocalypse, since Wikipedia also reports on the recent progress being made in quantum prime factorization:

"In 2001, Shor's algorithm was demonstrated by a group at IBM, who factored 15 into 3 times 5, using an NMR (nuclear magnetic resonance) implementation of a quantum computer with 7 qubits. After IBM's implementation, two independent groups implemented Shor's algorithm using photonic qubits, emphasizing that multi-qubit entanglement was observed when running the Shor's algorithm circuits. In 2012, the factorization of 15 was performed with solid-state qubits. Also, in 2012, the factorization of 21 was achieved, setting the record for the largest integer factored with Shor's algorithm. And three years ago, in 2019, an attempt was made to factor the number 35 using Shor's algorithm on an IBM Q System One, but the algorithm failed because of accumulating errors. Though larger numbers have been factored by quantum computers using other algorithms, these algorithms are similar to classical brute-force checking of factors, so unlike Shor's algorithm, they are not expected to ever perform better than classical factoring algorithms."

So, quantum computers have successfully factored the 4-bit value of 15, several times, and broke the record by factoring the 5-bit value of 21 using Shor's algorithm. But in an effort three years ago in 2019, couldn't quite make it to 6-bits to factor 35, which is binary 100011.

One of our listeners is a crypto-aware physicist who wrote after last week's podcast and he raised a couple of very good points which I want to share. We'll get to him in a minute, but I want to finish up on the asymmetric vs symmetric crypto question. Elsewhere, Wikipedia notes that: *"In contrast to the threat quantum computing poses to current public-key algorithms, most current symmetric cryptographic algorithms and hash functions are considered to be relatively secure against attacks by quantum computers. While Grover's quantum algorithm does speed up attacks against symmetric ciphers, doubling the key size can effectively block these attacks.*

Thus post-quantum symmetric cryptography does not need to differ significantly from current symmetric cryptography.”

And we already periodically double the lengths of our symmetric crypto keys and hashs as the speed of traditional computation begins to narrow their practical security margins.

Closing The Loop

Alim / @dutchphyscst
Hi Steve, Hope that all is well with you. As my weekly routine, I listened to your last podcast episode where you discussed Biden’s memorandum on Quantum Computer threats on classical cryptography. Being a physicist by education (PhD) and having practiced PKI-related IT work in the last 5-6 years, I wanted to make a few remarks on your comments.

It is not only RSA (based on difficulty of factorisation problem) but also ECC (based on discrete logarithm problems) which is vulnerable to Shor’s algorithm by a powerful enough quantum computer. Recently, Bruce Schneier also referred to an academic article where the authors discussed how much qubit capacity is required to achieve a reasonable attack on the Bitcoin blockchain. Indeed, the required number of physical qubits is tremendous (on the order of 10^6). On the other hand, we know how fast it went with the traditional silicon technology (remember Moore’s law).

I find Biden’s statement correct from some aspects:

- There is an attack called store-now decrypt later. Any stored information today, can be broken by a powerful quantum computer in future. Therefore, any confidential information that should stay confidential for a long period of time, should be protected against quantum computers today.

- Achieving crypto agility is very difficult. Mostly, cryptographic algorithms are embedded deep in the protocols and products. There are even cases that DES algorithm is still used in SIM cards and payment cards today. This makes the lifetime of such algorithms very long (i.e. 40 years or more). Therefore, Biden’s warning on the federal agencies and, thus, the industry is an early call for a very hectic and difficult transition.

Though, I also see some issues with Biden’s statement:

- NIST’s competition on post-quantum algorithms has not announced the winners yet and standard finalization is still a few years down the road. From this perspective, any organization to attempt to implement a post-quantum solution is a premature action. Things may still change.

- I sincerely hope that Biden’s administration did not implicitly try to pressure NIST to finalize the competition by this statement. Just a few months ago, a weakness is found and reported in one of the post-quantum digital signature algorithms (i.e. Rainbow). There should be no political pressure on such standardization activities and researchers should be left free in making their decision and given enough/proper time.

Needless to say, I definitely share your opinion that there are way more fundamental issues to be addressed urgently. However, I believe that quantum computer threats on cryptography is very serious and should be given enough attention due to slow adaptation of new cryptographic algorithms in billions of products/protocols etc.

With best regards,
Alim

Henrik Johnson / @henrikjohnson
Also just in the news:
https://www.engadget.com/ibm-wants-its-quantum-supercomputers-running-at-4000-plus-qubits-by-2025-110012129.html It doesn't seem like it is that far off until we start being in trouble given that there is no main stream asymmetrical crypto that is quantum safe.

[Name Redacted]
Good morning, Steve.

In Security Now, you've reported on a number of cybersecurity initiatives that the federal government has introduced this past year, including CISA's "Known Exploited Vulnerabilities Catalog", Congress' "Strengthening American Cybersecurity Act", and the White House's "Executive Order on Cybersecurity". What I haven't heard you mention are the TSA's two "Security Directive Pipeline" 2021-01 and 2021-02 memorandums. These are two successive directives, issued in response to the Colonial Pipeline compromise, that impose explicit cybersecurity requirements upon the midstream oil & gas pipeline industry.

One of the lesser-known regulatory mandates of the TSA (yes, that TSA) is the safety of interstate pipelines. I work in the midstream pipeline industry and these TSA directives have been the bane of my existence for the better part of a year. I'll reserve specific criticism, but will offer a recent Politico article which summarizes the situation nicely. Unfortunately, I'm not able to go into many particulars because the government, in its infinite wisdom, has marked the entire second directive (SD02) as Sensitive Security Information which prevents me from publicly divulging details. Suffice it to say, that yes, the government has instituted a cybersecurity standard that a segment of critical infrastructure must adhere to, but that can't be discussed except behind closed doors.

One tidbit that I am compelled to share is the role that CISA's "Known Exploited Vulnerabilities Catalog" plays. SD02 requires that pipeline operators patch vulnerabilities published in the Catalog within certain timeframes. Since you've mentioned the Catalog in several Security Now episodes, I wanted to call out that fact that this applies not just to government entities, but also to private pipeline companies. And yes, we are forced to review the list daily for new additions.

Thank you for all you do, and especially for a wonderful and informative weekly podcast.

[And thank you!... you've just contributed to making it more informative!]

Liam Lynch / @L2actual
Hi Steve,
I only listened to episode 869 the other day and I heard you and Leo again refer to the GDPR as being the cause of cookie notices. I kept meaning to contact you about this, as the only thing the GDPR did for cookie notices was to strengthen the consent requirement. Cookie notices have been around for a lot longer than the GDPR has been in force as they came from the ePrivacy directive and Rowena's short thread which I've shared with you here explains why they are so painful.
Love the show.
Regular listener ... Liam.

SpinRite

I wanted to officially note that the work on SpinRite's backend has officially finished. SpinRite's new hardware drivers are working without any exception that the group's extensive testing has revealed. In the case of two very old systems it was necessary to turn off the "UltraDMA" setting in the BIOS to obtain reliable transfers, but then everything worked perfectly. Throughout this work, SpinRite's new and much-improved benchmarking was used to exercise the backend drivers through the IO abstraction that I've talked about before. What that does is effectively isolate any backend devices from the front-end code. When SpinRite 7 adds native hardware USB drivers and then NVMe drivers, nothing else needs to change because the IO abstraction provides a uniform interface to the front-end code. The benchmark was the first client of that IO abstraction. The actual SpinRite machine, with its multiple switchable screens, grid display, DynaStat data recovery, detailed technical log and the rest will be the second and final client. That's where I now turn my focus. What we've just slogged and fought through has been by far the longest and toughest part, since it was where all of the hardware and machine dependency was. That's all now behind us.

Since the BIOS was historically SpinRite's IO abstraction, which is now gone, I have a lot of rewriting to do to support SpinRite's new abstraction. But it's not the sort of thing that will need constant interaction and tireless testing as the previous work on the backend did. I'm sure that the SpinRite testing gang will find things I've missed and will have ideas for improvements. But at this stage they're mostly going to be waiting for me rather than me waiting to learn from them how the latest test release turned out. :)

The New EU Surveillance State

The title of today's podcast might seem hyperbolic, but just wait ‘til you hear what the EU is proposing: The European Union's proposed new regulation will not only require scanning encrypted communications for child sexual abuse material content, but, believe it or not, actually reading all text messages with the goal of detecting any textual content that might be regarded as “grooming” a minor.

Those who haven't read far into the legislation quickly and correctly recognize that accomplishing any of this inherently, necessarily and unavoidably requires that some agency or entity scrutinizes all communications capable of conveying any graphical or textual material — in other words, all of the social platform messaging used by European Union citizens. And such scrutiny necessarily contravenes the well established goals, intents, and capabilities of end-to-end encryption. So, yeah, this would be the end of the true meaningful privacy enabled and facilitated by end-to-end encryption.

But reading some of the proposed legislation, as I did, one discovers that it also requires that this surveillance goes beyond the matching of previously known content hashes, to also include content that has not been previously seen. So, this would require either humans to view everything that everyone sends to anyone, and/or to train up machine vision and learning models to automate the identification of previously unknown child sexual abuse material.

Listen to what Johns Hopkins’ cryptographer, Matthew Green tweeted upon learning of this last week. Mathew Tweeted:

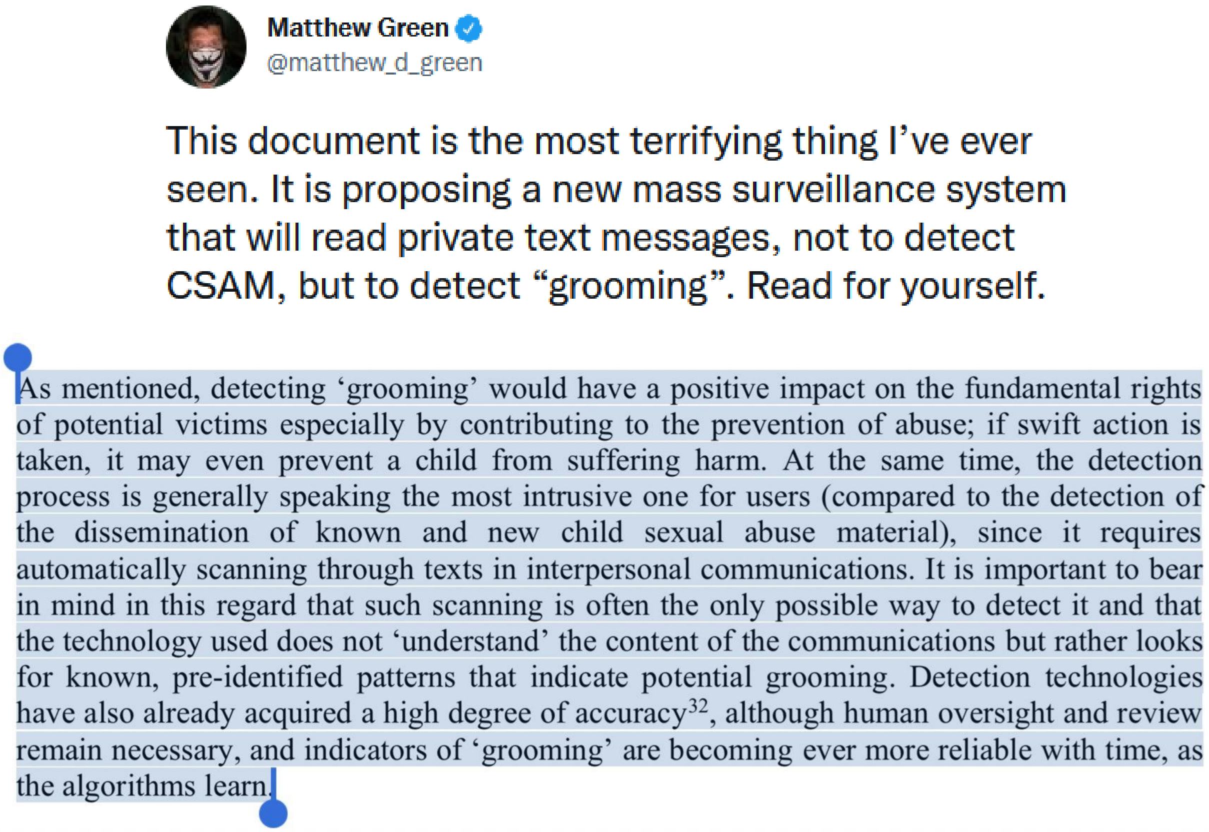

Matthew Green
@matthew_d_green

This document is the most terrifying thing I’ve ever seen. It is proposing a new mass surveillance system that will read private text messages, not to detect CSAM, but to detect “grooming”. Read for yourself.

As mentioned, detecting ‘grooming’ would have a positive impact on the fundamental rights of potential victims especially by contributing to the prevention of abuse; if swift action is taken, it may even prevent a child from suffering harm. At the same time, the detection process is generally speaking the most intrusive one for users (compared to the detection of the dissemination of known and new child sexual abuse material), since it requires automatically scanning through texts in interpersonal communications. It is important to bear in mind in this regard that such scanning is often the only possible way to detect it and that the technology used does not ‘understand’ the content of the communications but rather looks for known, pre-identified patterns that indicate potential grooming. Detection technologies have also already acquired a high degree of accuracy[32], although human oversight and review remain necessary, and indicators of ‘grooming’ are becoming ever more reliable with time, as the algorithms learn.

Matthew then issued a series of Tweets in a thread, writing:

> *Let me be clear what that means: to detect "grooming" is not simply searching for known CSAM. It isn't using AI to detect new CSAM, which is also on the table. It's running algorithms reading your actual text messages to figure out what you're saying, at scale. It is potentially going to do this on encrypted messages that should be private. It won't be good, and it won't be smart, and it will make mistakes. But what's terrifying is that once you open up "machines reading your text messages" for any purpose, there are no limits. Here is the document. It is long but worth reading, because it describes the most sophisticated mass surveillance machinery ever deployed outside of China and the USSR. Not an exaggeration.*

The link Matthew shared last week was from a leak of the official legislation which then appeared the next day. They are the same 135-page document. The title is "Proposal for a REGULATION OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL laying down rules to prevent and combat child sexual abuse."

https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=COM%3A2022%3A209%3AFIN&qid=1652451192472

As I said, the legislation is 135 pages. But its first two paragraphs set the stage and they're worth sharing:

> *The United Nations Convention on the Rights of the Child (UNCRC) and Article 24(2) of the Charter of Fundamental Rights of the European Union enshrine as rights the protection and care of children's best interests and well-being. In 2021, the United Nations Committee on the Rights of the Child underlined that these rights must be equally protected in the digital environment. The protection of children, both offline and online, is a Union priority.*
>
> *At least one in five children falls victim to sexual violence during childhood. A 2021 global study found that more than one in three respondents had been asked to do something sexually explicit online during their childhood, and over half had experienced a form of child sexual abuse online. Children with disabilities face an even higher risk of experiencing sexual violence: up to 68% of girls and 30% of boys with intellectual or developmental disabilities will be sexually abused before their 18th birthday.*
>
> *Child sexual abuse material is a product of the physical sexual abuse of children. Its detection and reporting is necessary to prevent its production and dissemination, and a vital means to identify and assist its victims. The pandemic has exposed children to a significantly higher degree of unwanted approaches online, including solicitation into child sexual abuse. Despite the fact that the sexual abuse and sexual exploitation of children and child sexual abuse materials are criminalised across the EU by the Child Sexual Abuse Directive 6 , adopted in 2011, it is clear that the EU is currently still failing to protect children from falling victim to child sexual abuse, and that the online dimension represents a particular challenge.*

One of the hardest lessons an ethical person learns and must come to terms with as they grow is that not all problems have workable solutions. And it doesn't matter at all how big the problem is nor how much we want there to be a good solution.

This entire problem can be broken down into two separate issues:

First, in order for some overviewing agency to obtain the raw data to be scrutinized, there can be no true and meaningful privacy between digitally communicating endpoints or individuals. That must end. Period. Everything needs to be visible and visited. And it's not sufficient to only surveil the devices used by minors because the original intent of detecting child sexual abuse material was to discover and apprehend those non-minors who were actively trading in such illegal content, thus curtailing the demand for the creation of more material. This means that all of everyone's social media messaging content must pass through surveillance filters.

Which brings us to the second issue: "What to do with that content once it's been obtained."

Matthew put it bluntly in one of his Tweets: *"It is going to do this on encrypted messages that should be private. It won't be good, and it won't be smart, and it will make mistakes."* Even if we agreed to voluntarily relinquish all of our privacy rights, it's not at all clear that the world has the technology to do what the EU's governing legislators want. It's easy for them to write a law stating what they want. But wanting it doesn't will it into existence... no matter how fervently and sincerely they want it. So the technology is going to miss things that it should catch and flag things that it should not. Humans will be required to examine the previously private photos that some image classifier believes to be salacious, and previously private text messages shared by consenting adults will be open to others' scrutiny.

And then there's the devil's advocate side which is also absolutely true and well-established:

Cryptography has already escaped. The algorithms which are able to unbreakably encipher plaintext are all public. So if the use of truly unbreakable end-to-end encryption is outlawed then only the outlaws will be using it. And, yes, that could be prevented, too. The next step in this escalation to doom would be for the communications carriers to refuse to transit any encrypted communications that they cannot themselves decrypt. That's possible. In which case we might as well just turn back the clock to the 1970's and give up because the Internet would no longer be useful for commerce.

Matthew also noted the other elephant in the room, Tweeting: *"But what's terrifying is that once you open up "machines reading your text messages" for any purpose, there are no limits."*

The distasteful issue of child sexual abuse is certainly very real. But it has been observed that it also serves as a convenient stalking horse for governments' much broader interests in monitoring and controlling speech of many other kinds. Much of such speech could be criminal. But much that might be of interest to censors would not be.

The EU's governors are wrong to want this legislation which can only be characterized as dangerous, wholly impractical, and impossible to implement as they hope. Pray that it dies.