Security Now! #752 - 02-04-20

The Little Red Wagon

This week on Security Now!

This week we examine the most recent flaw found in Intel's processors and what it means. We look at the continually moving target that is Windows 10, we consider the Free Software Foundation's suggestion that Microsoft open-source Windows 7 and the fact last months was apparently NOT the last update of Windows 7 for all non-ESU users. We look at the evolution of exploitation of the Remote Desktop Gateway flaw, Google's record breaking vulnerability bounty payouts, the return of Roskomnadzor, the size of fines, the question of who owns our biometrics, an update on AVAST/AVG spying, the future of 3rd-party A/V, a major milestone for the WireGuard VPN, and the wonderful little red wagon hack of the decade which titled this podcast.

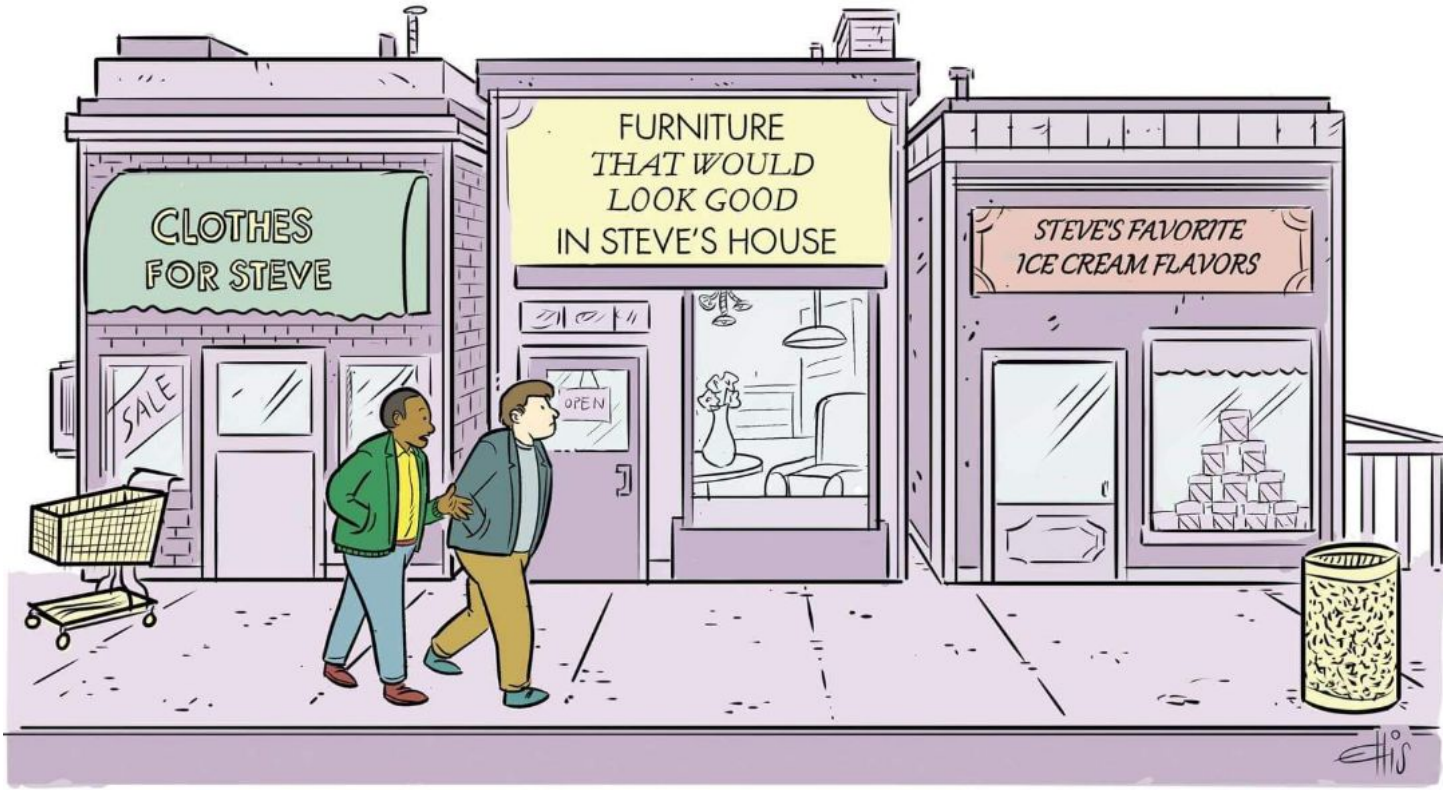

"Maybe you should disable your cookies, Steve."

Security News

L1D Eviction Sampling becomes "CacheOut"
So, Leo, last week at the beginning of the podcast you noted that news of another attack against Intel chips had surfaced. Since it bore the academic name "L1S Eviction Sampling" where "L1" would refer to the Level1 cache and "eviction" is the term used by caching systems when they must remove the typically least recently used (LRU) data to make room for new data to be cached. So from the name it sounded as though it must have been an attack that arranges to leverage the fact that the presence of caching would subtly and measurably alter the timing of cache-dependent code in such a way that sufficiently clever researchers could determine what had been evicted, and arrange to use that information to exfiltrate data across the processor's security boundaries. In other words, another headache for Intel.

Today, we have the more popular name for the new attack: "CacheOut" and, of course, it has a website and a logo: https://cacheoutattack.com/

The page's introduction says:

> CacheOut: Leaking Data on Intel CPUs via Cache Evictions
>
> We present CacheOut, a new speculative execution attack that is capable of leaking data from Intel CPUs across many security boundaries. We show that despite Intel's attempts to address previous generations of speculative execution attacks, CPUs are still vulnerable, allowing attackers to exploit these vulnerabilities to leak sensitive data.
>
> Moreover, unlike previous MDS (Microarchitectural Data Sampling) issues, we show in our work how an attacker can exploit the CPU's caching mechanisms to select what data to leak, as opposed to waiting for the data to be available. Finally, we empirically demonstrate that CacheOut can violate nearly every hardware-based security domain, leaking data from the OS kernel, co-resident virtual machines, and even SGX enclaves. "

The Researchers at the universities of Adelaide and Michigan demonstrated:

- The effectiveness of CacheOut in violating process isolation by recovering AES keys and plaintexts from an OpenSSL-based victim,
- Practical exploits for completely de-randomizing Linux's kernel ASLR, and for recovering secret stack canaries from the Linux kernel,
- How CacheOut effectively violates the isolation between two virtual machines running on the same physical core,
- How CacheOut could also be used to breach the confidentiality SGX guarantees by reading out the contents of a secure enclave,
- How some of the latest Meltdown-resistant Intel CPUs are still vulnerable, despite all of the most recent patches and mitigations.

Intel, who really used to enjoy their original job of printing money, once again responded to this latest annoyance:
https://software.intel.com/security-software-guidance/software-guidance/l1d-eviction-sampling

A speculative execution side channel variant known as L1D Eviction Sampling may allow the data value of some modified cache lines in the L1 data cache to be inferred under a specific set of complex conditions. L1D eviction sampling has been assigned CVE-2020-0549 with a CVSS of 6.5 CVSS:3.1/AV:L/AC:L/PR:L/UI:N/S:C/C:H/I:N/A:N.

On some processors under certain microarchitectural conditions, data from the most recently evicted modified L1 data cache (L1D) line may be propagated into an unused (invalid) L1D fill buffer. On processors affected by Microarchitectural Data Samping (MDS) or Transactional Asynchronous Abort (TAA), data from an L1D fill buffer may be inferred using one of these data sampling side channel methods. By combining these two behaviors together, it may be possible for a malicious actor to infer data values from modified cache lines that were previously evicted from the L1 data cache. This is called L1D eviction sampling.
Malicious software may be able to use L1D eviction sampling to infer modified cache line data written by previously run software, or modified cache line data written by software running on a sibling hyperthread on the same physical core.

Note that the term "modified" in "modified cache line data" refers to the fact that the data being evicted has been modified and that the cache is "write back" rather than "write through" so it needs to be written back to the external DRAM, which is what takes the time. If the data had not been modified then it could just be overwritten, which would take no time.

Unlike L1 Terminal Fault (L1TF), L1D eviction sampling doesn't potentially allow a malicious actor to select the physical address to probe.

Note that unless thread synchronization mitigations are applied, it may be possible for malicious software running on a sibling hyperthread to observe values loaded from or stored to memory on a physical core using the previously disclosed MDS or TAA methods.

As the list of processors affected by L1D eviction sampling are a subset of those affected by L1TF, systems affected by L1D eviction sampling may run software that already applies L1TF mitigations. Fully applying the L1TF mitigations for virtual machine managers (VMMs) ensures that the sensitive memory contents of the VMM or other virtual machines (VMs) will not be in the L1D cache when a possibly malicious VM executes. This helps prevent the malicious VM from attacking a VMM with L1D eviction sampling.
Mitigation

Intel expects to release microcode updates for affected processors which will mitigate the L1D eviction sampling issue. When the microcode update is released, software can discover if the microcode update contains the mitigation by reading the patch revision number and ensuring it matches or is greater than the corresponding revision number in the Affected Processors table.

The "CacheOut" page has an extensive FAQ for anyone who's interested. But the takeaway for most, if not all, of us is the same as it's been for the past two years:

- Yes, it's scary and it sounds bad, but...
- Even though it's real, it's another difficult to exploit fringe theoretical/academic problem.
- No known real world attacks have been detected in the wild. (But they would not likely be.)
- Intel has released another round of microcode patches.
- Linux can incorporate them easily into its boot. Windows 10 will be getting them.
- So basically: Nothing to see here, move along. Not much to worry about. But worth noting.

The researchers did note that it's unlikely for Antivirus products to detect and block CacheOut attacks, and since the exploit does not leave any traces in any traditional log file, it's also "very unlikely" to identify whether someone has exploited the flaw or not. And CacheOut flaw cannot be exploited remotely from a web browser and also does not affect AMD processors.

And speaking of microcode updates for Windows to load at boot time...
Last Thursday, Microsoft released updated Intel Microcode for Windows 10 1909, 1903, 1809, 1803, 1709, 1703 and 1607. However, the updates will not be automatically installed through Windows Update -- at least not currently. If wanted they must be downloaded and installed manually. Based upon the Windows build edition, the knowledgebase articles are:

- KB4497165: Intel microcode updates for Windows 1909 and 1903
- KB4494174: Intel microcode updates for Windows 1809
- KB4494451: Intel microcode updates for Windows 1803
- KB4494452: Intel microcode updates for Windows 1709
- KB4494453: Intel microcode updates for Windows 1703
- KB4494175: Intel microcode updates for Windows 1607
- KB4494454: Intel microcode updates for Windows 10

The firmware updates cover a huge range of Intel processors with added coverage for Denverton, Sandy Bridge, Valley View and Whiskey Lake architectures.

And by the way... I would never consider putting any of those things on my machines. When, or if, Microsoft ever decides it's necessary, I presume it'll be part of a monthly roll-up. But there is really no reason or need for an individual to do that. We all cherish the performance of the machines we have. Caching and speculative execution are all about performance. That's why they were created. So when those things are removed, so is performance.

Only one final version of Windows?
And, by the way... looking over those seven different knowledge base articles each for one or two separate editions of Windows... what-the-hell ever happened to there only being **ONE** Windows 10 from now on? That sounded like the **one** thing this new operating system had going for it. But I've been gaining some experience with it, and as far as I can see we now have more effectively separate and differing versions of Windows than we've ever had before! And it's not every six years. Now it's every six months!

Like most of us, whenever I hit a mystery with something I'm unfamiliar with, rather than struggling, I figure that if I'm having a problem a bunch of other people probably did too. So my "barrier to Google" is very low. I'll quickly "ask the Google" and see whether someone else has

asked and had their question answered. This has been a super-effective strategy with Windows 7. But I've been noticing that it's failing quite often with Windows 10. And the reason is that unlike Windows 7, there is no single Windows 10. Windows 10 is sort of a smear over time. Microsoft keeps messing with it, moving things around. Adding and removing things. Never leaving well enough alone. The result is that the amazing resource of other users having hit identical snags is not nearly as available or applicable as it once was. Because you cannot have identical snags when everyone is using a different version of Windows 10.

Windows 7 and the Free Software Foundation

And while we're talking about Window 10, it's worth mention that the Free Software Foundation, famously founded by Richard Stallman in 1985, has asked Microsoft to "Upcycle" Windows 7 (a term I've never heard before) by releasing it to the public. Uh huh. Oh, yeah… ***that's*** going to happen.

https://www.fsf.org/windows/upcycle-windows-7

Current signers:

We've reached the goal - but there's still time to show your support.

Sign the petition before February 5th to stay updated on the campaign.

On January 14th, Windows 7 reached its official "end-of-life," bringing an end to its updates as well as its ten years of poisoning education, invading privacy, and threatening user security. The end of Windows 7's lifecycle gives Microsoft the perfect opportunity to undo past wrongs, and to upcycle it instead.

We call on them to release it as free software, and give it to the community to study and improve. As there is already a precedent for releasing some core Windows utilities as free software, Microsoft has nothing to lose by liberating a version of their operating system that they themselves say has "reached its end."

To the executives at Microsoft:

- We demand that Windows 7 be released as free software. Its life doesn't have to end. Give it to the community to study, modify, and share.
- We urge you to respect the freedom and privacy of your users - not simply strongarm them into the newest Windows version.
- We want more proof that you really respect users and user freedom, and aren't just using those concepts as marketing when convenient.

We need your help to send Microsoft a strong message. We want 7,777 supporters to take a stand with us for freedom - not just for ourselves, but for future generations of computer users.

Please stand with us today, and sign below to show your support.

Last night they had obtained nearly twice their goal, but the entire thing seems much more like some sort of farcical stunt than anything that's serious. Yes, the source for the final version of MS-DOS was eventually released, but there's ZERO reason to imagine that Microsoft's reaction to this would be anything other than howling laughter. "We demand" ??? Really? Wow.

And speaking of Windows 7 updates

It turns out that the "supposed to be the last" January Patch Tuesday update for Windows 7 broke something that Microsoft cannot let lie. Believe it or not, it broke the desktop wallpaper stretch functionality which results in a black as night desktop when an image is being stretched to fit the desktop.

https://support.microsoft.com/en-us/help/4534310/windows-7-update-kb4534310

In their announcement of the January 14th, 2020 update Microsoft acknowledged the error, under "Known issues in this update" : After installing KB4534310, your desktop wallpaper might display as black when set to Stretch. Under "Workarounds":

To mitigate the issue, you can do one of the following:

- Set your custom image to an option other than Stretch, such as Fill, Fit, Tile, or Center.
- Choose a custom wallpaper that matches the resolution of your desktop.

We are working on a resolution and will provide an update in an upcoming release, which will be released to all customers running Windows 7 and Windows Server 2008 R2 SP1.

Their initial update indicated that this would ONLY be fixed for Windows 7 ESU customers. But the idea that the final update actually BROKE something that had always been working -- and something that's so obvious -- didn't sit well. So it appears there's going to be another update to Windows 7. Perhaps we'll get a few more security fixes at the same time.

RCE Exploit for Windows RDP Gateway Demoed by Researcher
We recently discussed the discovery of a Denial of Service attack on Microsoft's Remote Desktop Protocol Gateway service. And, once again, we see that when a sufficiently skilled hacker carefully examines nearly **any** sort of software flaw, it's often possible for that researcher to discover some way to manipulate the vulnerable machine to execute code of their choosing.

InfoGuard AG's penetration tester Luca Marcelli demonstrated a working remote code execution exploit. The exploit targets the RD Gateway on devices running Windows Server (2012, 2012 R2, 2016, and 2019). And it's being called "BlueGate."

Marcelli said that a blog post detailing how to achieve RCE with BlueGate will be forthcoming during the next few days, but that he wanted to be responsible and "wait a bit until people had enough time to patch before releasing this to the public."

To remind everyone, the RD Gateway allows admins to allow connections coming from the Internet to access Remote Desktop servers on internal networks only after proper authentication. So it's a sort of firewall proxy. But it has a vulnerability... Or rather two vulnerabilities which have been dubbed "BlueGate" They are both pre-authentication remote code execution and, now that they were patched in January's update cycle, Microsoft rates them both critical.

The trouble is, there are still nearly 20,000 unpatched and vulnerable RD Gateways accepting connections worldwide with 6,816 in the US alone. And since the use of RD Gateway is higher-end server platform only, it's likely that these are more than 6800 valuable networks.

Google more than doubles its own bug bounty record
During 2019 Google paid out $6.5 million in bug-bounty rewards, doubling the previous top annual total.

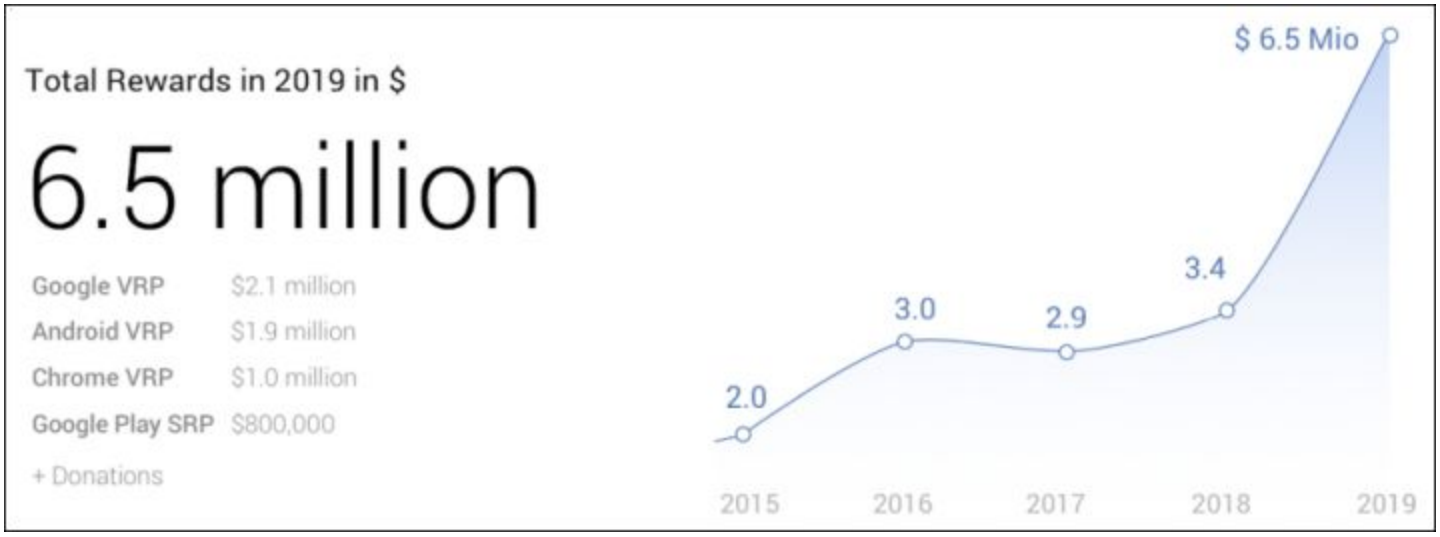

Google has been getting serious about bug bounties. Last year they launched their Developer Data Protection Reward Program aimed at uncovering data-abuse issues in Android apps, OAuth projects and Chrome extensions. They are looking for any apps that violate Google Play, Google

API and Google Chrome Web Store Extension privacy policies. Depending on the impact of the bug found, researchers may be rewarded up to $50,000 for a single report.

And also in 2019, Google tripled it top reward payouts for security flaws in Chrome from $5,000 to $15,000 – and doubled the maximum reward amount for high-quality reports from $15,000 to $30,000.

The Android Security Rewards program added additional exploit categories, and raised the top prize to $1 million for a full-chain, remote-code-execution exploit with persistence that compromises the Titan M secure element on Pixel devices. Recall that last May, Google recalled Bluetooth versions of the chip after discovering a vulnerability that would allow attackers within Bluetooth range to take control of the device.

And there's also the Google Play Security Reward Program which paid-out $650,000 in rewards during the second half of 2019, after it expanded its scope to any app (including third-party apps) with more than 100 million installs.

And recall the 50% bonuses we previously discussed. Google's security team added in their posting last Tuesday: "If you achieve [the top-reward Titan M Pixel exploit] on specific developer preview versions of Android, we're adding in a 50 percent bonus, making the top prize $1.5 million." The discovery of major problems in pre-release Android is worth much more to Google, which translates into much more for its discoverer.

And, at its 10th bug-bounty anniversary (Google began offering bounties in 2010), Google has paid a grand total of $21 million in rewards to date -- and a whopping $6.5 million of that in the most recent year.

The return of Roskomnadzor!

Roskomnadzor, which is, as we know, Russia's telecommunications watchdog, announced last Friday that it has instituted "administrative proceedings" (unquote) against Facebook and Twitter because of each company's refusal to move the data of Russian users to servers located inside the country's borders.

Roskomnadzor said: "These companies did not provide information on meeting the requirements for localizing the databases of Russian users of the corresponding social networks on servers located in the Russian Federation, as provided for in part 5 of Article 18 of the Law on Personal Data No. 152-03. Administrative proceedings have therefore been instituted on the grounds of an administrative offense in accordance with part 8 of article 13.11 of Administrative Code of the Russian Federation, which provides for an administrative fine in the amount of 1 million to 6 million rubles." ... Which is, wait for it... $16,000 to $94,000 USD.

As we've covered before, Facebook was previously threatened with a ban in September 2017 for the same reason. Twitter agreed to the demands of Russian officials at the time and proceeded to inform the Roskomnadzor that it was planning to move Russian users' data by mid-2018.

According to the Moscow Times, Roskomnadzor said Friday that a complaint will also be filed in Russian courts next week. And a new law signed last month by Vladimir Putin imposes higher

fines of up to 18 million rubles for repeat offenders. That's $280,500.

Meanwhile, last Wednesday the ProtonMail twitter account tweeted that: "The Russian government has blocked ProtonMail and ProtonVPN within Russia. We are reaching out to the appropriate authorities to get the block lifted as soon as possible."
— ProtonMail (@ProtonMail) January 29, 2020

This ban was prompted by Proton's refusal to register their services with Russian authorities — which was asked of all VPN providers operating in Russia. ProtonMail and ProtonVPN users are advised by the company to access the two services through Tor.

I wonder what's going on here? Does Russia want US social media companies operate inside Russia or not? Do they really care where their citizen's data are stored? If yes to either then why are years passing and fines so low relative to the sizes of the companies?

Facebook DID get fined, but not by Russia
Facebook just lost a class-action lawsuit it had been fighting for the past five years. And it will pay $550 million dollars in settlement of the suit. The lawsuit was brought against Facebook for scanning their user's faces in photos and offering tagging suggestions.

The plaintiffs claimed that the platform violated the strictest biometric privacy law in the land – Illinois's Biometric Information Privacy Act (BIPA) – due to its tag suggestions tool.

Facebook started using the tool in 2015 to automatically recognize people's faces in photos and suggest to their friends that they tag them. And it does this without users' permission and without telling them how long it would hang on to their biometrics. The lawsuit contends that Facebook squirreled away faceprints in what Facebook has claimed is the largest privately held database of facial recognition data in the world.

Last September, Facebook said it was ending tag suggestions in favor of the multi-purpose "face recognition" setting, which it made available to all users, along with an opt-out option.

The New York Times, in its reporting of the lawsuit outcome, referred to the $550 million hit as "a rounding error" for Facebook, which reported that revenue rose 25% to $21 billion in the fourth quarter, compared with a year earlier, while profit increased 7% to $7.3 billion.

Although this is a LOT of money (more than half a billion dollars) it could have been worse because Illinois' BIPA requires companies to get written permission before collecting a person's biometrics, whether they are fingerprints, facial scans or other identifying biological characteristics. It also gives Illinois residents the right to sue companies for up to $5,000 per violation, which could get pretty expensive. So, not surprisingly, Facebook fought this lawsuit tooth and nail. In 2016, it tried – and failed – to wriggle out of it by saying that its user agreement stipulates that California law would govern any disputes with the company. Besides, Facebook said in its motion, BIPA doesn't apply to Facebook's facial tagging suggestions for photos.

The judge's response was: nope, squared. Going by Illinois law was just fine, and it was always clear that BIPA would cover faceprints because it governs the use of all biometrics.

And speaking of BIPA...
We are clearly at a juncture, deciding who exactly owns our biometric data. 23 and Me just laid off 100 employees which was 14 percent of their workforce as consumer demand for its kits has weakened. The company said that a variety of factors, including privacy concerns, could have contributed to the slowing market. This followed a bunch of news coverage indicating that the company was responding to law enforcement queries. I was curious, so I looked up 23 and Me's own disclosure titled "How 23andMe Responds To Law Enforcement Requests For Customer Information"

https://customercare.23andme.com/hc/en-us/articles/212271048-How-23andMe-responds-to-law-enforcement-requests-for-customer-information

We work very hard to protect your information from unauthorized access from law enforcement. However, under certain circumstances, your information may be subject to disclosure pursuant to a judicial or other government subpoena, warrant or order, or in coordination with regulatory authorities. If such a situation arises, we have to comply with valid governmental requests and we will notify the affected individual(s) unless the legal request prevents us from doing so.

In other words, they must comply with a lawful court order and its likely accompanying gag order.

And apropos of this, a New York based facial recognition startup named "ClearView AI" has amassed a massive database containing more than three billion images scraped from employment sites, news sites, educational sites, and social networks including Facebook, YouTube, Twitter, Instagram and Venmo. And, since this data collection was done without the consent of the people whose images were collected, the company is now also being sued in a likely class action lawsuit. Since it uses BIPA, the complaint against Clearview AI was filed in Illinois.

The New York Times published an exposé about how Clearview has been quietly selling access to faceprints and facial recognition software to law enforcement agencies across the US, claiming that it can identify a person based on a single photo, revealing their real name and far more. From the New York Times:

> The tool could identify activists at a protest or an attractive stranger on the subway, revealing not just their names but where they lived, what they did and whom they knew.

Clearview told the Times that more than 600 law enforcement agencies have started using Clearview in the past year, and it's sold the technology to a handful of companies for security purposes. Clearview declined to provide a list of its customers.

Eric Goldman, co-director of the High Tech Law Institute at Santa Clara University, told the newspaper that the "weaponization possibilities" of such a tool are "endless." Imagine a rogue law enforcement officer who wants to stalk potential romantic partners, or a foreign government using this to dig up secrets about people to blackmail them or throw them in jail.

The New York Times headline about ClearView AI suggested that the company "might end privacy as we know it." From the report:

> Even if Clearview doesn't make its app publicly available, a copycat company might, now that the taboo is broken. Searching someone by face could become as easy as Googling a name. Strangers would be able to listen in on sensitive conversations, take photos of the participants and know personal secrets. Someone walking down the street would be immediately identifiable – and his or her home address would be only a few clicks away. It would herald the end of public anonymity.

The complaint claims that Clearview's technology gravely threatens civil liberties: "Constitutional limits on the ability of the police to demand identification without reasonable suspicion, for instance, mean little if officers can determine with certainty a person's identity, social connections, and all sorts of other personal details based on the visibility of his face alone."

The lawsuit claims that Clearview isn't just selling this technology to law enforcement, it's also allegedly sold its database to private entities including banks and retail loss prevention specialists and has developed ways to implant its technology in wearable glasses that private individuals could use.

Welcome to our Sci-Fi future. It's here sooner than we expected.

I have nothing to hide and when I'm out in public I'm not in disguise. But since I'm not a high visibility celebrity, there's a sort of assumption that people who I don't know, don't know me, just as I don't know them. But it would change things if we all had an app like a "super Shazzam" that we could point at anyone and instantly obtain their name, address, full biographical background and links to their various social media accounts. That would change the world.

Jumpshot missed the hoop

So, we all remember that AVAST and AVG were both found to be massively tracking their users behind their backs. The cover story for the tracking was that the user's browsing was being sent back to the mothership so that the URLs could be checked for safety... even though other browser-based safety checkers offer the same service locally and without sending the user's URL click stream to a central database. The tracking turned out to be quite extensive, essentially sending back ALL of a user's actions, such as changing tabs and scrolling the page, to provide detailed monitoring of their users. The data was anonymized, but the researcher who discovered it (the well known author of Ad Block Plus) felt that deanonymizing the data wouldn't be difficult.

But that wasn't the point or the concern. What was creepy was that AVAST had purchased a data analytics firm that was apparently reselling this data to whomever wanted to troll through it.

Mozilla reacted to the news by immediately ousting AVAST and AVGas browser extensions from the Mozilla repository. We subsequently heard that AVAST/AVG had dramatically reduced the amount of data being collected, and I wondered aloud on this podcast how their prior business

model could be sustained if so much less data was collected.
We, the other shoe has dropped, and it turns out that the answer is: It cannot be. Avast is winding down its subsidiary named "Jumpshot" following this investigation into the sale of their users' data to third parties that may pose a risk to their privacy. Last Thursday, the antivirus vendor said the unit will no longer have access to user information harvested from users of Avast products and services will eventually be fully terminated.

https://www.jumpshot.com/

Jumpshot has ceased operations. Thank you.

Jumpshot was purchased in 2013 and began life under Avast as a PC cleanup tool. But then, five years ago the subsidiary's business shifted to data analytics and its focus pivoted to marketing intelligence based on the analysis of online consumer spending patterns and purchases. Jumpshot reportedly had access to information from over 100 million devices.

A joint investigation conducted by Motherboard and PCMag, just published, revealed that information scraped by Avast from users and handed over to Jumpshot is linked to individuals through a unique ID in an effort to anonymize them -- but it is possible to pick apart data strings to de-anonymize users and reveal their identity, tracing their online footprint, browsing habits, and purchases.

In a blog post, Avast's CEO said that the recent news about Jumpshot "has hurt the feelings of many of you and rightfully raised a number of questions," and as chief executive, he feels "personally responsible" for the turmoil.

And speaking of 3rd-party antivirus vendors...
To no one's surprise, every single anti-virus vendor has formally announced that they will continue offering and fully supporting their existing A/V solutions for Windows 7 for at least the next two years, through 2022. And we can realistically expect them to support Windows 7 until the actual installed base shrinks to the point that it is no longer profitable for them to offer A/V for that platform.

And, given that Windows Defender just suddenly stopped protecting all Windows 7 users, this clearly represents a significant marketing opportunity for the A/V industry.

https://www.zdnet.com/article/all-major-antivirus-vendors-will-continue-to-support-windows-7-post-eol/

Do we have any recommendations for our listeners moving forward?

An Update on the WireGuard VPN in the Linux kernel
The lean-coded, fast, modern, and secure WireGuard VPN protocol has made it into the Linux kernel as Linus Torvalds merged it into his source tree for version 5.6.

The next Linux kernel is expected to be released in just a few months. And in the list of kernel changes, the WireGuard VPN was the first thing on the list.

So what's behind all the excitement?

The WireGuard protocol, and its implementation, is a project from security researcher and kernel developer Jason Donenfeld, who created it as an alternative to IPsec and OpenVPN. In its current form, WireGuard has about 4,000 lines of code versus the 100,000 lines of code that OpenVPN has been dragging forward and it doesn't require OpenSSL.

Compared to current OPenVPN's "kitchen sink something for everyone" options, WireGuard relies upon a small set of carefully chosen modern cryptographic primitives that are stronger, perform better, have been highly scrutinized by the cryptographic community:

- ChaCha20 for symmetric encryption, authenticated with Poly1305, using RFC7539's AEAD construction
- Curve25519 for ECDH (elliptic-curve Diffie-Hellman) key agreement
- BLAKE2s for hashing and keyed hashing, described in RFC7693
- SipHash24 for hashtable keys
- HKDF for key derivation, as described in RFC5869

WireGuard provides perfect forward secrecy, protection against denial-of-service, brute-force attacks, key impersonation, and replay attacks, as well as support for an additional layer of symmetric-key cryptography to offer some post-quantum resistance.

The lean selection of crypto primitives shows that it has deliberately dropped all of the unnecessary choices for encryption, key encryption, and hashing algorithms. This increases interoperability since there's less to mismatch while also minimizing the risks from the support of obsolete crypto. Moreover, WireGuard is also faster because it lives in the kernel space, it's easier to audit for security vulnerabilities, and it's simple to configure and deploy.

https://www.wireguard.com/papers/wireguard.pdf

And it's not just for Linux. It is also fully cross-platform with implementations for Windows, macOS, BSD, iOS, and Android.

And some VPN service providers (Mullvad, AzireVPN, IVPN, VPN.ac, and TorGuard) already offer WireGuard servers. And allow me to reiterate that while I was in Gothenberg Sweden for the SQRL OWASP presentation there, I happened to meet and break bread with the founder of Mullvad VPN. Remember how surprised I was when I signed up for sync.com and the "Send me an eMail for password recovery" checkbox was UN-checked by default? I was so impressed by that. That's the way I felt during dinner as the guys from Mullvad casually talked, described their company and shared anecdotes about the unforeseen consequences of deliberately wanting to know NOTHING about their customers -- including their IP addresses which are never logged.

In this week's Best Hack of the New Decade...
(Even if the new decade doesn't actually start until next year!)

Why would an ingenious German performance artist place 99 active cell phones into a little red wagon and walk them around the streets of Berlin?? Because it brilliantly confused and spoofed Google Maps into believing that there must be massive traffic congestion, since so many "cars" were all moving so slowly through a region of the city!

https://youtu.be/k5eL_al_m7Q

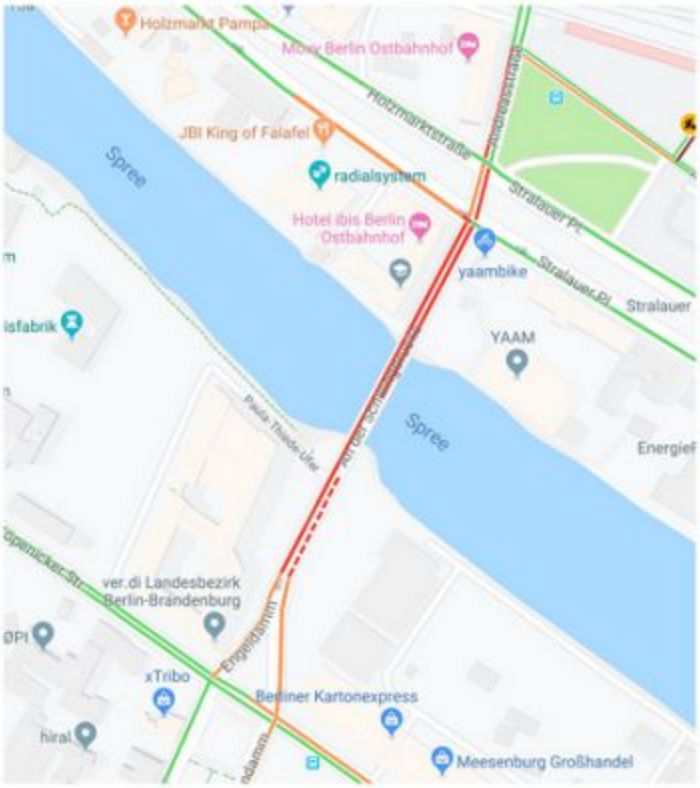

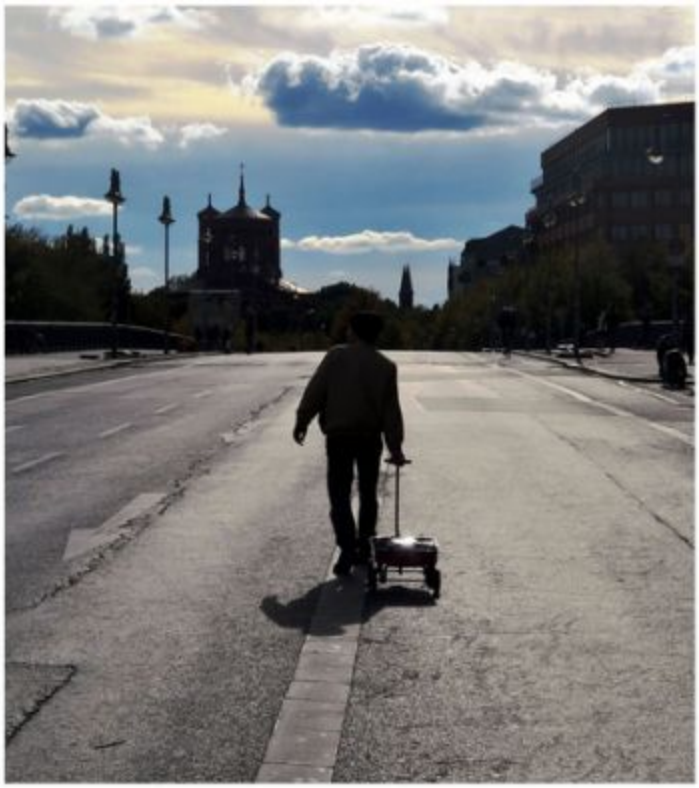

It's unclear what time of day this was, but the video and the photos show that he's completely alone on the streets… presumably because Google Maps freaked out drivers who depend upon it for their directions and chose to find another way to their destination.