Security Now! #1013 - 02-18-25

Chrome Web Store is a mess

This week on Security Now!

US lawmakers respond to the UK's outrageous demand about Apple's encryption. What, exactly, is a "backdoor", and can a "backdoor" **NOT** be secret? Highlights from last week's Windows' Patch Tuesday. A look into RansumHub: The latest king of the Ransomware hill. "TOAD": Telephone-Oriented Attack Delivery. The state of Texas -versus- DeepSeek. Disabling Apple's "Restricted Mode". Where did I put that $800 million in Bitcoin? A Sci-Fi author update. And a deep dive into the misoperation of Chrome's critically important Web Extension Store.

Lest there be any doubt...

Security News

"US lawmakers respond to the UK's Apple encryption backdoor request"

Last Thursday, Engadget's gave their updated coverage of the UK decryption order that headline, with the sub-head "Sen. Ron Wyden and Rep. Andy Biggs said the order is 'effectively a foreign cyberattack waged through political means.' " Engadget wrote:

The UK's shockingly intrusive order for Apple to create a backdoor into users' encrypted iCloud data doesn't only affect Brits; it could be used to access the private data of any Apple account holder in the world, including Americans. Less than a week after security experts sounded the alarm on the report, the U.S. Congress is trying to do something about it.

The Washington Post reported on Thursday that, in a rare show of modern Capitol Hill bipartisanship, Sen. Ron Wyden (D-OR) and Rep. Andy Biggs (R-AZ) wrote to the new National Intelligence Director Tulsi Gabbard, asking her to take measures to thwart the UK's surveillance order — including limiting cooperation and intelligence sharing if the country refuses to comply.

Biggs and Wyden wrote: "If Apple is forced to build a backdoor in its products, that backdoor will end up in Americans' phones, tablets and computers, undermining the security of Americans' data, as well as of the countless federal, state and local government agencies that entrust sensitive data to Apple products. The US government must not permit what is effectively a foreign cyberattack waged through political means."

The pair told Gabbard that if the UK doesn't retract its order, she should "reevaluate US-UK cybersecurity arrangements and programs as well as US intelligence sharing with the UK." Wyden sits on the Senate Intelligence Committee, and Biggs is on the House Judiciary Committee and chairs the Subcommittee on Crime and Federal Government Surveillance.

Wyden began circulating a draft bill that, if passed, could at least make the process harder for UK authorities. The proposed modification to the 2018 CLOUD Act would make information requests to US-based companies by foreign entities more onerous by requiring them to first obtain a judge's order in their home country. In addition, it would forbid other countries (like, say... the UK) from demanding changes in encryption protocols to the products or services of companies in the US. Request challenges would also be given jurisdiction in US rather than foreign courts.

The UK order, first reported by The Washington Post, requires Apple to create a backdoor into its Advanced Data Protection, a feature introduced in iOS 16.2 in 2022. Advanced Data Protection applies end-to-end encryption to many types of iCloud data, including device backups, Messages content, notes and photos, making them inaccessible even to Apple. The order demands a blanket ability to access a user's fully encrypted data whenever and wherever the target is located.

The order was issued under the UK's Investigatory Powers Act 2016, known (not so affectionately) as the "Snooper's Charter," which expanded the electronic surveillance powers of British intelligence agencies and law enforcement. It would be a criminal offense for Apple to publicly confirm receiving the order, so the company hasn't commented on the matter. Security experts warn that implementing this backdoor would needlessly expose anyone with an Apple Account to foreign spying, hackers and adversarial countries.

Apple received a draft of the order last year when UK officials debated the changes. In a written submission protesting them, the company said the planned order "could be used to force a company like Apple, that would never build a back door into its products, to publicly withdraw critical security features from the UK market." The company can appeal the notice but cannot use the appeal to delay compliance.

Ciaran Martin, former chief executive of the UK's National Cyber Security Center, told The Washington Post: "Most experts in the democratic world agree that what the UK is proposing would weaken digital security for everyone, not just in the UK but worldwide."

I wanted to take a moment to focus upon the use of the term "Backdoor." Unfortunately, its original meaning is being lost and stretched through reuse for other purposes. The term was used liberally throughout the original Washington Post article, and also in Engadget's own reporting, which I've just shared. Engadget quoted the two U.S. Senators saying:

"If Apple is forced to build a backdoor in its products, that backdoor will end up in Americans' phones, tablets and computers, undermining the security of Americans' data, as well as of the countless federal, state and local government agencies that entrust sensitive data to Apple products."

In the past, I've pedantically objected to the use of the term "Backdoor" in these cases and I'm going to take this opportunity to be at least as pedantic about this again today. I have previously suggested that what's being asked for is a locked, yet unlockable, front door. I suppose the trouble is that this stuff can be confusing for those who don't inhabit the security space for a living. The term "Backdoor" sounds bad and "bad" is often the way someone wants it to sound.

So what's wrong with using the term "Backdoor"? My problem is that words need to have and to hold onto their meaning and "Backdoor" already has an extremely specific and exact meaning. It was originally used to describe any sort of security measure bypass, and it was *definitely* meant to be a *secret*. Period. A backdoor is, by definition, a secret. So the UK cannot possibly mandate the inclusion of a "Backdoor" into anything, because anything mandated could never be a secret. The UK could certainly mandate that Apple have some means for complying with their demands for a users' data, and if that data was initially encrypted for the user's privacy, then Apple would need to have some means for decrypting it in order to comply with the UK's demand. But nothing about that suggests the use of any sort of "Backdoor" and, in fact, from where we are now, Apple would need to deliberately design-in a new "front door" for which only they possess the key. Apple clearly objects to doing this and for that I salute them. As has been previously mentioned, Google has supported full end-to-end, device-to-device encryption of cloud stored data from Android 9 the "Pie" edition, and in this case Pie referred to dessert rather than to Pre-Internet Encryption, even though that's what it offered.

So if we should not refer to designed-in decryption capabilities as "Backdoors" what should they be called? The problem is, the security industry doesn't have any sufficiently pithy and engaging term for this. So "Backdoor" it is, for better or for worse, even though that isn't at all what anyone is asking for – whether or not they know it.

Anyway, I did say I was going to be pedantic about this, and I'm sure I haven't disappointed on that count. Every time I see the term "backdoor" – which has a very specific – meaning being used as a generic term for obtaining otherwise-inaccessible information I think to myself "yeah, but a backdoor is not what it is."

Patch Tuesday Review
Compared with last month's massive batch of software fixes, February's updates were mild. They addressed a mere 63 flaws and eliminated a pair of less severe, though still actively-exploited, 0-day flaws in Windows. Of those 63 flaws, three were rated Critical, 57 were deemed to be merely Important, one was Moderate, and the last two were rated as low severity. In addition to those 63, Microsoft also separately resolved 23 flaws in their Chromium-based Edge browser.

The two resolved 0-days had CVSS scores of 7.1 and 7.8. The CVSS 7.1 was an elevation of privilege in Windows Storage. Microsoft's alert said: *"An attacker would only be able to delete targeted files on a system. This vulnerability does not allow disclosure of any confidential information, but could allow an attacker to delete data that could include data that results in the service being unavailable."* However, Mike Walters, the president and co-founder of Action1, noted that the vulnerability could be chained with other flaws to escalate privileges and perform follow-on actions that can complicate recovery efforts and allow threat actors to cover up their tracks by deleting crucial forensic artifacts.

The second 0-day having the higher CVSS of 7.8 also created an elevation of privilege vulnerability, this time in Windows Ancillary Function Driver for WinSock. WinSock is short for Windows Sockets and is part of the operating system's networking subsystem. Due to the fact that the ADF.SYS driver is in the kernel, the successful exploitation of this vulnerability would allow an attacker to obtain SYSTEM privileges.

A similar flaw in AFD.SYS was disclosed by Gen Digital last August after they found that it had been weaponized by North Korea's Lazarus Group. And a year ago, in February 2024, Microsoft plugged a Windows kernel privilege escalation flaw affecting the AppLocker driver (appid.sys) that was also being actively exploited by the same group. These attack chains stand out because they do rely upon the Bring Your Own Vulnerable Driver (BYOVD) approach. Instead they take advantage of comparatively rare security flaws in native Windows drivers to eliminate the need to introduce vulnerable drivers into their targets.

It's not known whether the abuse of last month's 0-day is also linked to the Lazarus Group. CISA has added both the flaws to its Known Exploited Vulnerabilities (KEV) catalog. Their presence in CISA's KEV catalog requires federal agencies to apply patches by March 4, 2025.

The most severe of the flaws addressed by Microsoft in this month's update is CVE-2025-21198 carrying a CVSS of 9.0. It's a remote code execution (RCE) vulnerability in the High Performance Compute (HPC) Pack. Microsoft documented that *"An attacker could exploit this vulnerability by sending a specially crafted HTTPS request to the targeted head node or Linux compute node granting them the ability to perform RCE on other clusters or nodes connected to the targeted head node."* Although this is bad, it wasn't known to be abused at the time of its patching. Now, however, the vulnerability has become known after the fact, so bad guys could potentially reverse engineer the vulnerability from the update and start attacking any exposed and still-vulnerable Windows networks.

There's also an 8.1 CVSS which affects Windows LDAP, its Lightweight Directory Access Protocol. The flaw allows an attacker to send a specially crafted request and to execute arbitrary code. Since that's really not good, the LDAP flaw would have a higher CVSS if the vulnerability was not so difficult to exploit. The flaw is a race condition which reduces the chances to a craps shoot having a low chance of success. Still, Ben McCarthy, the lead cybersecurity engineer at Immersive Labs, said: "Given that LDAP is integral to Active Directory, which underpins authentication and access control in enterprise environments, a compromise could lead to lateral movement, privilege escalation, and widespread network breaches."

And speaking of authentication, there was also a CVSS 6.5 NTLM v2 hash disclosure vulnerability which, if successfully exploited, would permit an attacker to authenticate as the targeted user.

So we don't have any "sky is falling" updates, but, as usual, updating as soon as practical would be a good idea.

RansomHub is 2024's Top Ransomware Group, Hitting 600+ Organizations Globally

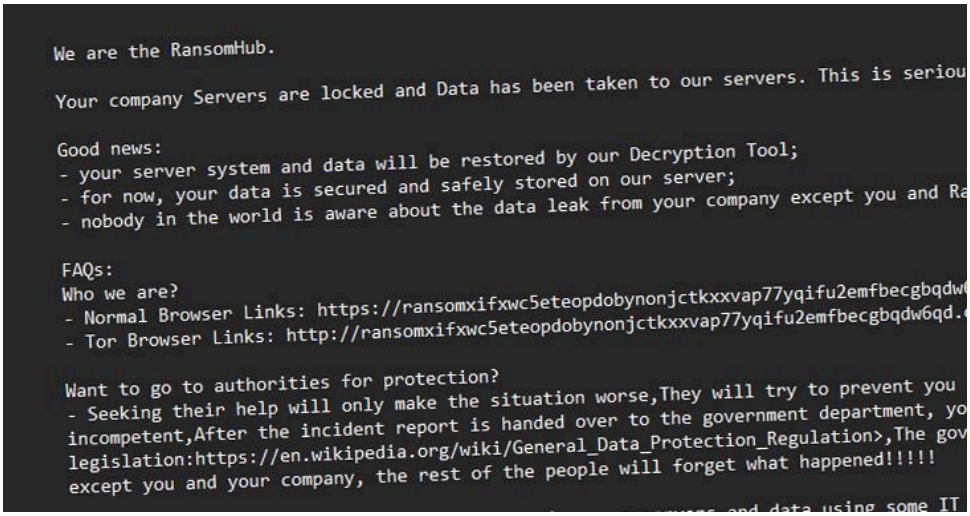

A new and quite effective ransomware-as-a-service (RaaS) group, who call themselves "RansomHub", rose in prominence to become last year's #1 ransomware perpetrator after compromising the networks and data of more than 600 organizations worldwide.

The RansomHub bad guys have been observed leveraging now-patched security flaws in Microsoft Active Directory and the Netlogon protocol to escalate privileges and gain unauthorized access to a victim network's domain controller as part of their post-compromise strategy.

Analysts at Group-IB write in a report published last week that "RansomHub has targeted over 600 organizations globally, spanning sectors including healthcare, finance, government, and critical infrastructure. This has firmly established them as the most active ransomware group in 2024."

The group first surfaced exactly one year ago, in February 2024, after acquiring the source code associated with the now-defunct Knight (formerly known as Cyclops) RaaS gang from the RAMP cybercrime forum. Five months later, an updated version of the locker was advertised on the illicit marketplace with capabilities to remotely encrypt data via the Simple File Transfer Protocol (SFTP). The group's updated malware comes in multiple variants that are capable of encrypting files on Windows, VMware ESXi, and SFTP servers. RansomHub has also been observed actively recruiting affiliates from LockBit and BlackCat groups as part of a partnership program, indicating an attempt to capitalize on the law enforcement actions targeting its rivals.

In the incident analyzed by Group-IB, RansomHub unsuccessfully attempted to exploit a critical flaw impacting Palo Alto Networks PAN-OS devices (CVE-2024-3400) using a publicly available proof-of-concept (PoC), before ultimately breaching the victim network by means of a brute-force attack against the VPN service. The Group-IB researchers said: *"This successful brute force attack used an enriched dictionary of over 5,000 usernames and passwords. The attacker finally eventually gained entry through a default account frequently used in data backup solutions,*

which then allowed them to breach the network perimeter." The initial access was then used to carry out the ransomware attack, with both data encryption and exfiltration occurring within 24 hours of the compromise. The attack weaponized two known security flaws in Active Directory (CVE-2021-42278 aka noPac) and the Netlogon protocol (CVE-2020-1472 aka ZeroLogon) which allowed the attackers to seize control of the domain controller and conduct lateral movement within and across the network.

The researchers said that *"The exploitation of these vulnerabilities enabled the attacker to gain full privileged access to the domain controller, which is the nerve center of a Microsoft Windows-based infrastructure. Following the completion of the exfiltration operations, the attacker prepared the environment for the final phase of the attack. The attacker operated to render all company data saved on the various Network Attached Storage systems completely unreadable and inaccessible, as well as impermissible to restore, with the aim of forcing the victim to pay the ransom to get their data back."*

The researchers added: *"The origins of the RansomHub group, its offensive operations, and its overlapping characteristics with other groups confirm the existence of a still-active cybercrime ecosystem. This environment thrives on the sharing, reusing, and rebranding of tools and source code, fueling a robust underground market where high-profile victims, infamous groups, and substantial sums of money play central roles."*

Ransomware-as-a-Service affiliates are incentivized with an 80% share of ransom proceeds.

After originally being saturated in ransomware stores, I've been actively avoiding them since there hasn't been much new to report. Law enforcement has successfully tracked down and stomped out many of the larger and highest-profile groups. But, exactly as was predicted, any members who managed to escape law enforcement sweeps, or those who were more peripheral, changed groups, moved and merged into others, or formed new groups. The problem is, as we saw during last week's detailed look into attacks on K-12 school systems, there's just too much money potentially waiting to be collected from insurers for bad guys to ignore the chance to get some of it. So ransomware, in one form or another, promises to remain a cybercrime staple for the foreseeable future.

Android's "anti-TOAD" (Telephone-Oriented Attack Delivery)

Here's something that I didn't realize was "a thing" until I learned that Google was beta testing its prevention. There's a class of attack using the acronym "TOAD" which stands for Telephone-Oriented Attack Delivery. This forthcoming feature for Android 16 blocks fraudsters from sideloading apps during calls. When I read that I thought "sideloading apps during calls?" That's a thing that happens? The Hacker News explains:

> *Google is working on a new security feature for Android that blocks device owners from changing sensitive settings when a phone call is in progress. Specifically, new in-call anti-scammer protections include preventing users from turning on settings to install apps from unknown sources and granting accessibility access. The development was first reported by Android Authority.*

So, apparently, scammers are instructing unwitting users to do things during phone calls, such as, I suppose, when calling a fake technical support hotline for assistance. The Hacker News continues:

Users who attempt to do so during phone calls are served the message: "Scammers often request this type of action during phone call conversations, so it's blocked to protect you. If you are being guided to take this action by someone you don't know, it might be a scam." Furthermore, it blocks users from giving an app access to accessibility over the course of a phone call.

The feature is currently live in Android 16 Beta 2, which was released last week. With this latest addition, the idea is to introduce more friction to a tactic that has been commonly abused by malicious actors to deliver malware. Dubbed telephone-oriented attack delivery, these approaches involve sending SMS messages to prospective targets and instructing them to call a number by inducing a false sense of urgency.

Last year, NCC Group and Finland's National Cyber Security Centre (NCSC-FI) disclosed that cybercriminals were distributing dropper apps using a combination of SMS messages (To initiate scam calls) followed by phone calls to trick users into installing malware such as Vultr.

The development comes after Google expanded restricted settings to cover more permission categories in order to prevent sideloaded apps from accessing sensitive data. Google has also rolled out the ability to automatically block sideloading of potentially unsafe apps in markets like Brazil, Hong Kong, India, Kenya, Nigeria, Philippines, Singapore, South Africa, Thailand, and Vietnam to tackle fraud.

This seems like a useful feature and it's the sort of thing that our phones could easily do. Just notice that a call is in progress and raise the bar on potentially dangerous behavior that might be the result of instructions being received during that active telephone conversation. And this should serve as a reminder of just how effective social engineering attacks remain. Most people have no idea how any of this stuff works. So when a knowledgeable voice at the other end of the phone explains how to fix some made-up problem, many people will follow along, especially when this is the same thing that authentic helpers also do.

Texas -vs- DeepSeek

Under the heading "Because... why not?" we have the news, reported by The Record, that Texas is investigating "DeepSeek" which, you know, comes from China. What did they do wrong? They embarrassed the U.S. by making a better AI. So we've decided that they probably violated the state's data privacy laws, and we need to find out. In their reporting, The Record wrote:

Attorney General Ken Paxton's office also has requested relevant documents from Google and Apple, seeking their "analysis" of the inexpensive and open source DeepSeek app and asking what documentation they required from DeepSeek before they made the app publicly available for download on their app stores.

Paxton said in a statement: "DeepSeek appears to be no more than a proxy for the CCP [Chinese Communist Party] to undermine American AI dominance and steal the data of our citizens. That's why I'm announcing a thorough investigation and calling on Google and Apple to cooperate immediately by providing all relevant documents related to the DeepSeek app."

In other words, "Their AI is better than ours, and we can't have any of that! So we're going to investigate them in order to hopefully find some evidence of misbehavior. The Record wrote:

DeepSeek, Google and Apple did not immediately respond to requests for comment.

On January 28, Paxton banned DeepSeek's use on all devices owned by members of his staff due to security concerns and what a press release from his office called "the company's blatant allegiance to the CCP, including its willingness to censor any information critical of the Chinese government."

This week, New York state and Virginia both blocked the use of DeepSeek on government devices, and on Monday, Reps. Josh Gottheimer (D-NJ) and Darin LaHood (R-IL) introduced a bipartisan bill that would ban federal workers from using DeepSeek on government devices.

Sadly, anti-Chinese technology backlash has become predictable. With DeepSeek just being the latest example. Since it's exceedingly difficult to prove that China is **not** using their DeepSeek app to monitor the questions, behavior and who knows what else of U.S. citizens, it appears that we're inevitably heading into a world of increasing mistrut — a technology cold war — where everyone is only going to be trusting the hardware, software and firmware produced by their country and their close allies. And even close allies are having trouble, as we're seeing with the emerging stand-off between the UK and Apple.

That this was where we appeared to be headed has also been clear for years. As tensions between the U.S. and both China and Russia have been gradually mounting, everyone listening to this podcast has heard me wonder on many occasions how it is that China and Russia were still using Microsoft's Windows – an operating system that could so easily be hiding pro-Western capabilities. As we know, both countries have felt similarly and are now working to remove Windows from their critical enterprises and industries. And that's a feat that's more easy ordered than accomplished.

Disabling restricted mode on Apple devices

I wanted to note that eight days ago Apple announced that they had updated all of their operating systems to fix a bug that they said may have been used in "extremely sophisticated attacks against specific targeted individuals."

Back when it was introduced we covered the introduction of "Restricted Mode". It further locks down Apple devices where it's enabled. On the one hand, it makes those devices less fun to use because they can then do less. But in return it also makes them far less easy to compromise. I strongly endorsed the addition of this option since we still haven't figured out how to make highly complex products 100% secure and bulletproof.

The flaw that was fixed would have, and presumably did for a time, allow sophisticated attackers to employ it in an attack chain. It's role in the chain was to disable restricted mode on a locked device. The vulnerability, as described, could have been used to enable unlocking technology similar to that of Cellebrite's products, which allow snoopers to break into devices when they have physical access to them. Apple's restricted mode helps with this by proactively blocking data access to iPhones and iPads when they have been locked for more than an hour. It's very clever.

The vulnerability in Apple's iOS and iPadOS affects iPhone XS and later, iPad Pro 13-inch, iPad Pro 12.9-inch 3rd generation and later, iPad Pro 11-inch 1st generation and later, iPad Air 3rd generation and later, iPad 7th generation and later, and iPad mini 5th generation and later, Apple said.

$800 Million in Bitcoins
In other news, we have James Howells, that poor guy who lost his hard drive containing the only copy of the 51-character private key he needs to unlock his cryptocurrency wallet. This is an issue because that wallet contains 8,000 bitcoins, currently valued at around $800 million dollars with a single Bitcoin now worth $100,000. Ouch. That's gotta hurt. James is certain that the drive was mistakenly thrown out with the trash and is now lurking somewhere in trash landfill in Newport City, Wales.

Last month, James lost a court battle with the Newport City Council in Wales, which may have been his last shot at excavating the dump since soon after, the city council revealed that it would be closing the landfill and building a large solar farm on the site. In a press release the city council said that the solar farm would help the city replace its fleet of diesel garbage trucks with EVs to help the city transform a site that otherwise would not be suitable for other uses into a meaningful effort to reduce the city's carbon footprint. I'm sure such stories abound.

Sci-Fi Update

I've been waiting to gain sufficient experience with a new-to-me Sci-Fi author before mentioning my recent science fiction reading enjoyment. I took ChatGPT up on its advice about other authors who were similar to those whose novels I have previously enjoyed – often more than once. As we recall, ChatGPT not only produced a list of recommendations, but among those were others of my favorites that I had never mentioned. Still being cautiously suspicious of AI, we wondered whether ChatGPT might have previously ingested my published Sci-Fi reading list, or the transcripts of this podcast.

But in any event, I obtained a handful of new author recommendations. Since I had seen Neal Asher's name around a lot, I purchased a copy of "Gridlinked". And I do mean purchased. It wasn't free as part of the Kindle Unlimited plan that everything else I've been reading recently has been. But given that inflation has jacked the price of a 5-shot Starbucks Venti Latte all the way to $9.50, — ouch! — paying $7 for a novel that will give me more than a week of true enjoyment works for me — so long as the novel is good. And so far I feel that I've easily obtained my money's worth.

I started with "Gridlinked" because it was Asher's early work and I prefer to start at the beginning. But if the critics on Reddit know what they're talking about, this 5-novel series, of which Gridlinked is the first, pales in comparison to Asher's later work. Someone who finished Gridlinked asked whether the other four in the series were worth reading. Someone replied:

> *I think he was finding his feet in the Polity universe with Gridlinked - his following works are miles ahead, keep at it, you won't be disappointed.*

Well, that sounds great because I'm already not disappointed. I've mentioned that I seem to be quite sensitive to an author's ability to write. It's not just the plot and the characters for me. They need to be able to express themselves. That previous book I tried to read "Artifact", where the alien starship was named the "Ziggawatt" lost me before the end of that first sentence.

The Goodreads site describes "Gridlinked" by writing:

> *Gridlinked is a science fiction adventure in the classic, fast-paced, action-packed tradition of Harry Harrison and Poul Anderson, with a dash of cyberpunk and a splash of Ian Fleming*

added to spice the mix.

Ian Cormac is a legendary Earth Central Security agent, the James Bond of a wealthy future where "runcibles" (matter transmitters controlled by AIs) allow interstellar travel in an eye blink throughout the settled worlds of the Polity. Unfortunately Cormac is nearly burnt out, having been "gridlinked" to the AI net for so long that his humanity has begun to drain away. He has to take the cold-turkey cure and shake his addiction to having his brain on the net.

It's a bit freaky that Neal Asher wrote about 'net addiction and the tendency to lose one's humanity through being over-connected, back in 2001 — 24 years ago when this book was first published. I won't say much more other than that I'm now 67% through the second of this 5-book series and I am really enjoying them. Asher's "Polity" universe is run by dispassionate AI's because humans cannot be trusted to wield such power. Within "the Polity" life is sweet and orderly with no crime and everyone has something interesting to do. So in that way it's reminiscent of life within Star Trek's Federation of planets. Of course, there are those who chafe under the bit of authority, and who prefer the freedom of anarchy. So, plenty of adventure, war and opportunity can be found out on the fringe beyond the control of the Polity.

Mostly, Neal Asher can write and I think he's a terrific storyteller. I will definitely keep paying $7 for the next three books, and given the Reddit comments about Neal's follow-on works, I'm going to be glad that I took ChayGPT up on its suggestions for similar authors.

Listener Feedback

Bob McNaughton

This might be obvious, but surely if you configure DNS over TLS in your browser, you will miss out on the caching performed by any of the more local DNS resolvers, such as the one in your router? Wouldn't it be better to use DNS over TLS in the router, thus hiding your DNS queries from your ISP, but getting the advantage of the cached lookups other people on the same LAN have performed?

Bob is 100% correct, of course. In all of our discussion, I had not mentioned that if a user configures their local web browser to use any form of encrypted DNS service – which seems to be the way things are evolving – some loss of local caching, for example by the local router if it does DNS caching, would be lost.

The flip side of this is that the emerging DNS Benchmark continues so show that once a TCP and TLS connection have been negotiated and brought up, the individual DNS lookups being offered by the Internet's major providers are actually being resolved FASTER by them than by, for example, even my own ISP's much more local resolvers.

As we noted last week, this might be due to the fact that encrypted DNS servers are still lightly loaded because the use of DNS over TLS or DNS over HTTPS is still the exception more than the rule. But I'm going to be very interested to learn what everyone discovers once the Benchmark can be more widely used.

Chrome Web Store is a mess

"Chrome Web Store is a mess" is the exact title someone who should know gave to a recent blog posting of his a few weeks ago. Wladimir Palant's posting caught my eye both due to his pedigree and due to the importance of his message. Anyone who's been following this podcast for more than a few years could probably reduce the number of major security trouble sources to a high single digit. And among those most important would be the security of web browser extensions.

Extensions to the basic functionality of our web browsers have been with us since nearly the beginning. And 20 years ago, back when there was much less to do on the Internet, the security of an add-on was much less crucially important. But every year since then, more and more of our lives have moved online. This has meant that the overall security and privacy offered by the web browsers we use to interact with the Internet has become increasingly important. And no one who has listened to more than a couple of this podcast's episodes could entertain any doubt that, disheartening though it may be, the world is apparently filled with an astonishing number of total strangers who would hurt us without a second thought to obtain any advantage.

Several times in recent weeks I've focused our attention upon the security and privacy issues surrounding web browser add-ons – sadly, there are many. So when I saw that Wladimir Palant had taken the time to push back a bit from the entrails of specific add-ons to survey the larger picture, I knew that was something I wanted to share.

Earlier, I mentioned Wladimir's pedigree, but his name may not ring any bells right off. So here's how Wladimir explains on his blog site. He writes:

> *My name is Wladimir Palant and I'm mostly blogging about security topics these days. You will often see me taking apart browser extensions because I've been developing those myself since 2003. One particularly well-known project of mine is* ***Adblock Plus*** *which I originally developed. Eventually, I co-founded eyeo, a company to take care of this project. I'm still developing the browser extension PfP: Pain-free Passwords while my other extensions have become obsolete over time.*
>
> *My writing is meant to help people learn. So I aim to provide information on both how vulnerabilities can be found and how these can be prevented in your own code. I won't merely discuss security issues but also try to draw generic conclusions from those and give recommendations.*
>
> *Despite researching security topics since at least 2007 I still do it as a hobby rather than my job. I experimented with earning money via bug bounty programs, which resulted in acceptable income. However, other aspects eventually turned me away from bug bounties. In particular, I want to write about my research and don't want to be prevented from it by a company taking years to fix an issue.*

In other words, he was becoming annoyed that after finding and reporting some problem – and being paid for his responsible disclosure – the bug bounty agreement would require that he never reveal anything about the problem until it had been fixed. This differs from unpaid security researchers who are able to set 90-day "fix it before we publish it" deadlines. So Wladimir was becoming annoyed that bugs were being purchased and he was being effectively gagged when he wanted to be able to document the problems and use them as illustrative teaching examples.

In any event, here's a highly technical developer who created one of the earliest and most popular privacy extensions and who has been at this for more than 22 years. When this guy titles his blog posting "Chrome Web Store is a mess" I want to understand why he thinks so.

Wladimir wrote:

Let's make one thing clear first: I'm not singling out Google's handling of problematic and malicious browser extensions because it is worse than Microsoft's for example. No, Microsoft is probably even worse but I never bothered finding out. That's because Microsoft Edge doesn't matter, its market share is too small. Google Chrome on the other hand is used by around 90% of the users world-wide, and one would expect Google to take their responsibility to protect its users very seriously, right? After all, browser extensions are one selling point of Google Chrome, so certainly Google would make sure they are safe?

Unfortunately, my experience reporting numerous malicious or otherwise problematic browser extensions speaks otherwise. Google appears to take the "least effort required" approach towards moderating Chrome Web Store. Their attempts to automate all things moderation do little to deter malicious actors, all while creating considerable issues for authors of legitimate add-ons. Even when reports reach Google's human moderation team, the actions taken are inconsistent, and Google generally shies away from taking decisive actions against established businesses.

As a result, for a decade my recommendation for Chrome users has been to stay away from Chrome Web Store if possible. Whenever extensions are absolutely necessary, it should be known who is developing them, why, and how the development is being funded. Just installing some extension from Chrome Web Store, including those recommended by Google or "featured," is very likely to result in your browsing data being sold or worse.

Google employees will certainly disagree with me. Sadly, much of it is organizational blindness. I am certain that Google meant well and that they did many innovative things to make it all work. But looking at it from the outside, it's the result that matters. And for the end users the result is a huge (and rather dangerous) mess.

Some recent examples

Five years ago I discovered that Avast browser extensions were spying on their users.

Remember that we covered that at the time. It was this guy who made the discovery, which may be why his name is at least somewhat familiar. He continues:

Mozilla and Opera disabled the extension listings immediately after I reported it to them. Google on the other hand took two weeks where they supposedly discussed their policies internally. The result of that discussion was eventually their "no surprises" policy which says:

> *Building and maintaining user trust in the Chrome Web Store is paramount, which means we set a high bar for developer transparency. All functionalities of extensions should be clearly disclosed to the user, with no surprises. This means we will remove extensions which appear to deceive or mislead users, enable dishonest behavior, or utilize clickbaity functionality to artificially grow their distribution.*

So when dishonest behavior from extensions is reported today, Google should act immediately and decisively, right? Let's take a look at two examples that came up in the past few months.

In October I wrote about the refoorest extension deceiving its users. I could conclusively prove that Colibri Hero, the company behind refoorest, deceives their users on the number of trees they supposedly plant, incentivizing users into installing with empty promises. In fact, there is strong indication that the company never even donated for planting trees beyond a rather modest one-time donation.

Google got my report and dealt with it. What kind of action did they take? That's a very good question that Google won't answer. But refoorest is still available from Chrome Web Store, it is still "featured" and it still advertises the very same completely made up numbers of trees they supposedly planted. Google even advertises for the extension, listing it in the "Editors' Picks extensions" collection, probably the reason why it gained some users since my report. So much for being honest. For comparison: refoorest used to be available from Firefox Add-ons as well but was already removed when I started my investigation. Opera removed the extension from their add-on store within hours of my report.

But maybe that issue wasn't serious enough? After all, there is no harm done to users if the company is simply pocketing the money they claim to spend on a good cause. So also in October I wrote about the Karma extension spying on users. Users are not being notified about their browsing data being collected and sold, except for a note buried in their privacy policy. Certainly, that's identical to the Avast case mentioned before and the extension needs to be taken down to protect users?

Again, Google got my report and dealt with it. And again I fail to see any result of their action. The Karma extension remains available on Chrome Web Store unchanged, it will still notify their server about every web page its users visit. The users still aren't informed about this. Yet their Chrome Web Store page continues to claim "This developer declares that your data is not being sold to third parties, outside of the approved use cases," a statement contradicted by the extension's privacy policy. The extension appears to have lost its "Featured" badge at some point but now it is back.

Note: Of course Karma isn't the only data broker that Google tolerates in Chrome Web Store. I published a guest article today by a researcher who didn't want to disclose their identity, explaining their experience with BIScience Ltd., a company misleading millions of extension users to collect and sell their browsing data. This post also explains how Google's "approved use cases" effectively allow pretty much any abuse of users' data.

Neither refoorest nor Karma were isolated instances. Both recruited or purchased other browser extensions as well. These other browser extensions were turned outright malicious, with stealth functionality to perform affiliate fraud and/or collect users' browsing history. Google's reaction was very inconsistent here. While most extensions affiliated with Karma were removed from Chrome Web Store, the extension with the highest user numbers (and performing affiliate fraud without telling their users) was allowed to remain for some reason.

With refoorest, most affiliate extensions were removed or stopped using their Impact Hero SDK. Yet when I checked more than two months after my report two extensions from my original list still appeared to include that hidden affiliate fraud functionality and I found seven new ones that Google apparently didn't notice.

The reporting process

Now you may be wondering: if I reported these issues, why do I have to guess what Google did in response to my reports? Keeping developers who report in the dark is Google's official policy:

> Hello Developer,
>
> Thank you again for reporting these items. Our team is looking into the items and will take action accordingly. Please refer to the possible enforcement actions and note that we are unable to comment on the status of individual items.
>
> Thank you for your contributions to the extensions ecosystem.
>
> Sincerely,
>
> Chrome Web Store Developer Support

This is the same response I received in November after pointing out the inconsistent treatment of the extensions. A month later the state of affairs was still that some malicious extensions got removed while other extensions with identical functionality were available for users to install, and I have no idea why that is. I've heard before that Google employees aren't allowed to discuss enforcement actions, and your guess is as good as mine as to whom this policy is supposed to protect.

Supposedly, the idea of not commenting on policy enforcement actions is hiding the internal decision making from bad actors, so that they don't know how to game the process. If that's the theory however, it isn't working. In this particular case the bad actors got some feedback, be it through their extensions being removed or due to the adjustments demanded by Google. It's only me, the reporter of these issues, who is left guessing.

But, and this is a positive development, I've received a confirmation that both these reports are being worked on. This is more than I usually get from Google which is: silence. And typically also no visible action either, at least until reports start circulating in media publications forcing Google to act on it.

But let's take a step back and ask ourselves: how does one report Chrome Web Store policy violations? Given how much Google emphasizes their policies, there should be an obvious way. In fact, there is a support document for reporting issues. And when I started asking around, even Google employees would direct me to it.

If you find something in the Chrome Web Store that violates the Chrome Web Store Terms of Service, or trademark or copyright infringement, let us know.

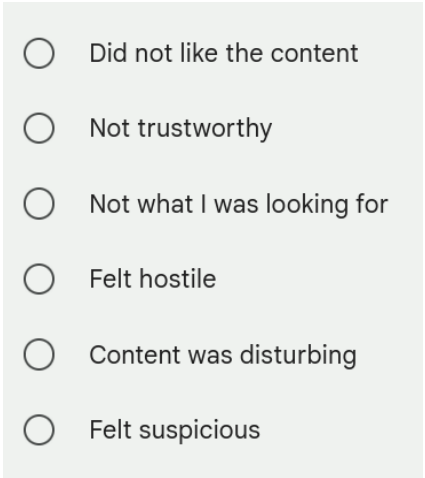

This doesn't really seem like the place to report policy violations. Even "Felt suspicious" isn't right for an issue you can prove. And, unsurprisingly, after choosing this option Google just responds with:

> *Your abuse report has been submitted successfully.*

No way to provide any details. No asking for my contact details in case they have questions. No context whatsoever, merely "felt suspicious." This is probably fed to some algorithm somewhere which might result in… what actually? Judging by malicious extensions where users have been vocally complaining, often for years: nothing whatsoever. This isn't the way.

Well, there is another option listed in the document:

> *If you think an item in the Chrome Web Store violates a copyright or trademark, fill out this form.*

Yes, Google seems to care about copyright and trademark violations, but a policy violation is not that. If we try the form nevertheless, it gives us a promising selection:

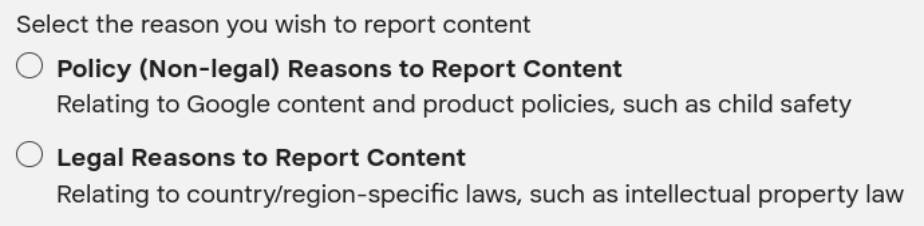

Finally! Yes, policy reasons are exactly what we are after, let's click that. And there comes another choice:

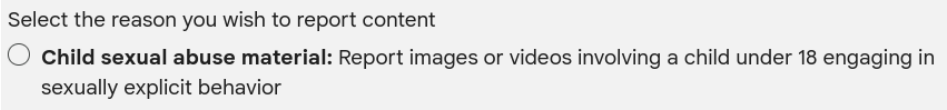

That's really the only option offered. And I have questions. At the very least those are: in what jurisdiction is child sexual abuse material a non-legal reason to report content? And: since when is that the only policy that Chrome Web Store has?

We can go back and try "Legal Reasons to Report Content" of course but the options available are really legal issues: intellectual properties, court orders or violations of hate speech law. So this is another dead end.

It took me a lot of asking around to learn that the real (and well-hidden) way to report Chrome Web Store policy violations is Chrome Web Store One Stop Support.

I mean: I get it that Google must be getting lots of non-sense reports. And they probably want to limit that flood somehow. But making legitimate reports almost impossible can't really be the way.

In 2019 Google launched the Developer Data Protection Reward Program (DDPRP) meant to address privacy violations in Chrome extensions. Its participation conditions were rather narrow for my taste, pretty much no issue would qualify for the program. But at least it was a reliable way to report issues which might even get forwarded internally. Unfortunately, Google discontinued this program in August 2024.

It's not that I am very convinced of DDPRP's performance. I've used that program twice. First time I reported Keepa's data exfiltration. DDPRP paid me an award for the report but, from what I could tell, allowed the extension to continue unchanged. The second report was about the malicious PDF Toolbox extension. The report was deemed out of scope for the program but forwarded internally. The extension was then removed quickly, but that might have been due to the media coverage. The benefit of the program was that it was a documented way of reaching a human being at Google that would look at a problematic extension. Now it's gone.

Chrome Web Store and their spam issue

In theory, there should be no spam on Chrome Web Store. The policy is quite clear on that:

> *We don't allow any developer, related developer accounts, or their affiliates to submit multiple extensions that provide duplicate experiences or functionality on the Chrome Web Store.*

Unfortunately, this policy's enforcement is lax at best. Back in June 2023 I wrote about a malicious cluster of Chrome extensions. I listed 108 extensions belonging to this cluster, pointing out their spamming in particular:

> *Well, 13 almost identical video downloaders, 9 almost identical volume boosters, 9 almost identical translation extensions, 5 almost identical screen recorders are definitely not providing value.*

I've also documented the outright malicious extensions in this cluster, pointing out that other extensions are likely to turn malicious as well once they have sufficient users. And how did Google respond? The malicious extensions have been removed, yes. But other than that, 96 extensions from my original list remained active in January 2025, and there were of course more extensions that my original report didn't list. For whatever reason, Google chose not to enforce their anti-spam policy against them.

And that's merely one example. My most recent blog post documented 920 extensions using tricks to spam Chrome Web Store, most of them belonging to a few large extension clusters. As it turned out, Google was made aware of this particular trick a year before my blog post already. And again, for some reason Google chose not to act.

What about extension reviews? Can they be trusted?

When you search for extensions in Chrome Web Store, many results will likely come from one of the spam clusters. But the choice to install a particular extension is typically based on reviews. Can at least these reviews be trusted? On the topic of moderation of reviews Google says:

> *Google doesn't verify the authenticity of reviews and ratings, but reviews that violate our terms of service will be removed.*

And the important part in the terms of service is:

> *Your reviews should reflect the experience you've had with the content or service you're reviewing. Do not post fake or inaccurate reviews, the same review multiple times, reviews for the same content from multiple accounts, reviews to mislead other users or manipulate the rating, or reviews on behalf of others. Do not misrepresent your identity or your affiliation to the content you're reviewing.*

Now you may be wondering how well these rules are being enforced. The obviously fake review on the Karma extension is still there, three months after being posted. Not that it matters, with their continuous stream of incoming five star reviews.

A month ago I reported an extension to Google that, despite having merely 10,000 users, received 19 five star reviews on a single day in September – and only a single (negative) review since then. I pointed out that it is a consistent pattern across all extensions of this account. For example, another extension (with only 30 users) received 9 five star reviews on the same day. It really doesn't get any more obvious than that. Yet all these reviews are still online.

I have a lot to say in reaction to what Wladimir is observing and reporting. But I'm holding that until he's finished. Still, I want to note that the automated clean-up of clearly bogus reviews would be trivial to implement. Wladimir is made suspicious when an extension with 30 users acquires nine 5-star reviews, all on the same day. Right. One wonders whether they were all posted from different accounts at the same IP address? Google would know that. But even if not, the fraudulent pattern is glaringly obvious.

And remember that it's more than likely that this conduct is also reflected in the operation of the extension itself. Someone who's unwilling to honestly earn a reputation for their extension is more likely to have ulterior motives. So if Google were to automate extension review clean-up, which, again, would be trivial for them to do, they would be reducing the damage being done through the fraudulent over-promotion of less savory extensions.

Because no trivial clean-up is happening, we need to wonder whether review spamming may be something Google doesn't mind, even if it's clearly hurting Chrome's users.

And it isn't only fake reviews. The refoorest extension incentivizes reviews which violates Google's anti-spam policy which says:

> *Developers must not attempt to manipulate the placement of any extensions in the Chrome Web Store. This includes, but is not limited to, inflating product ratings, reviews, or install counts by illegitimate means, such as fraudulent or incentivized downloads, reviews and ratings.*

It has been three months, and they are still allowed to continue. The extension gets a massive amount of overwhelmingly positive reviews, users get their fake trees, everybody is happy. Well, other than the people trying to make sense of these meaningless reviews.

With reviews being so easy to game, it looks like lots of extensions are doing it. Sometimes it shows as a clearly inflated review count, sometimes it's the overwhelmingly positive or meaningless content. At this point, any user ratings with the average above 4 stars likely have been messed with.

The "featured" extensions

But at least the "Featured" badge is meaningful, right? It certainly sounds like somebody at Google reviewed the extension and considered it worthy of carrying the badge. At least Google's announcement indeed suggests a manual review:

> *Chrome team members manually evaluate each extension before it receives the badge, paying special attention to the following:*
> 1. *Adherence to Chrome Web Store's best practices guidelines, including providing an enjoyable and intuitive experience, using the latest platform APIs and respecting the privacy of end-users.*
> 2. *A store listing page that is clear and helpful for users, with quality images and a detailed description.*

Yet looking through 920 spammy extensions I reported recently, most of them carry the "Featured" badge. Yes, even the endless copies of video downloaders, volume boosters, AI assistants, translators and such. If there is an actual manual review of these extensions as Google claims, it cannot really be thorough.

To provide a more tangible example, Chrome Web Store currently has Blaze VPN, Safum VPN and Snap VPN extensions all carrying the "Featured" badge.

These extensions (along with Ishaan VPN which has barely any users) belong to the PDF Toolbox cluster which produced malicious extensions in the past. A cursory code inspection reveals that ***all four are identical*** *and are, in fact, clones of Nucleus VPN which was removed from Chrome Web Store in 2021. And they also don't even work, no connections succeed. The extension not working is something users of Nucleus VPN complained about, which the extension compensated for by loading it with fake reviews.*

And, again, all of these carry the "Featured extension" badge. So it looks like the main criteria for awarding the "Featured" badge are the things which can be easily verified automatically: user count, Manifest V3, claims to respect privacy (not even the privacy policy, merely that the right checkbox was checked), and a Chrome Web Store listing with all the necessary promotional images. Given how many such extensions are plainly broken, the requirements on the user interface and general extension quality don't seem to be too high. And providing unique functionality definitely isn't on the list of criteria.

In other words: if you are a Chrome user, the "Featured" badge is completely meaningless.

It is no guarantee that the extension is not malicious, not even an indication. In fact, authors of malicious extensions will invest some extra effort to get this badge. That's because the website algorithm seems to weigh the badge considerably towards the extension's ranking.

How did Google get into this mess?

Google Chrome first introduced browser extensions in 2011. At that point the dominant browser extensions ecosystem was Mozilla's, having been around for 12 years already. Mozilla's extensions suffered from a number of issues that Chrome developers noticed:

Essentially unrestricted extension privileges necessitated very thorough reviews before extensions could be published on Mozilla Add-ons website. And since these extension code reviews largely relied on volunteers, they often took a long time, with publication delays being very frustrating to add-on developers.

Note that I was an extension reviewer on Mozilla Add-ons myself between 2015 and 2017.

Google Chrome was meant to address all these issues. It pioneered sandboxed extensions which allowed limiting extension privileges. And Chrome Web Store focused on automated reviews from the very start, relying on heuristics to detect problematic behavior in extensions, so that manual reviews would only be necessary occasionally and after the extension was already published. Eventually, market pressure forced Mozilla to adopt largely the same approaches.

Google's over-reliance on automated tools caused issues from the very start, and it certainly didn't get any better with the increased popularity of the browser. Mozilla accumulated a set of rules to make manual reviews possible. For example, all code should be contained in the extension, so no downloading of extension code from web servers. Also, reviewers had to be provided with an unobfuscated and unminified version of the source code. Google didn't consider any of this necessary for their automated review systems. So when automated review failed, manual review was often very hard or even impossible.

It's only with the recent introduction of Manifest V3 that Chrome finally prohibits remotely hosted code. And it took until 2018 to prohibit code obfuscation, while Google's reviewers still have to reverse minification for manual reviews. Mind you, we are talking about policies that were already long established at Mozilla when Google entered the market in 2011.

And extension sandboxing, while without doubt useful, didn't really solve the issue of malicious extensions. I already wrote about one issue back in 2016:

> *The problem is: **useful** extensions will usually request "give me the keys to the kingdom" permission. So these permissions always need to be granted.*

Essentially, this renders permission prompts useless. Users cannot possibly tell whether an extension has valid reasons to request extensive privileges. So legitimate extensions have to constantly deal with users who are confused about why the extension needs to "read and change all your data on all websites." Eventually, users become desensitized and trained to simply accept such prompts without thinking twice.

And then malicious add-ons come along, requesting extensive privileges under a pretense. Monetization companies put out guides for extension developers on how they can request more privileges for their extensions while fending off complaints from users and Google alike. There is a lot of this going on in the Chrome Web Store, and Manifest V3 is unable to change anything about it.

So what we have now is:

1. *Automated review tools that malicious actors willing to invest some effort can work around.*

2. *Lots of extensions with the potential for doing considerable damage, yet little way of telling which ones have good reasons for that and which ones abuse their privileges.*

3. *Manual reviews being very expensive and unreliable thanks to historical decisions.*

4. *Massively inflated extension count due to unchecked spam.*

Those last two ("Manual reviews being very expensive and unreliable thanks to historical decisions" and "Massively inflated extension count due to unchecked spam") further trap Google in the "it needs to be automated" mindset. Yet adding more automated layers isn't going to solve the issue when there are companies which can put a hundred employees on devising new tricks to avoid triggering detection. Yes, hundreds of employees because malicious extensions make a lot of money and are big business.

What could Google do?

If Google were interested in making Chrome Web Store a safer place, I don't think there is a way around investing considerable (manual) effort into cleaning up the place. Taking down a single extension won't really hurt the malicious actors, they have hundreds of other extensions in the pipeline. Tracing the relationships between extensions on the other hand and taking down the entire cluster – that would change things.

As the saying goes, the best time to do this was a decade ago. The second best time is right now, when Chrome Web Store with its somewhat less than 150,000 extensions is certainly

large but not yet large enough to make manual investigations impossible. Besides, there is probably little point in investigating abandoned extensions – those whose latest release is more than two years ago – which make up almost 60% of the Chrome Web Store.

But so far, Google's actions have been entirely reactive, typically limited to extensions which already caused considerable damage. I don't know whether they actually want to stay on top of this. From the business point of view there is probably little reason for that. After all, Google Chrome no longer has to compete for market share, having essentially won against the competition. Even with Chrome extensions not being usable, Chrome will likely stay the dominant browser.

Okay.

As we often observe on this podcast, it's certainly useful to tell someone to be careful when they may be considering some action that might have negative consequences for them. But at least for me, if I'm told not to do something, in order to really accept that I want to understand why – I want to understand exactly **why** something would be bad for me. I think that's why I grew up to respect my father. He was an explainer. So I suppose I come by that honestly.

His explaining approach always made sense to me because, armed with a full understanding, no one needs to tell me anything about what to do or not to do, since I'm able to judge that for myself.

So in the case of Google Chrome Web Store extensions, I'm not going to tell anyone not to download and install extensions they feel that they need. Rather, everyone who's reached this point in today's podcast is now fully equipped to judge for themselves whether anything that's there may be worth their time.

It would be great if Google were able to function as a reliable curator of the 135,000 Chrome Web Store extensions that are currently available for download. We now absolutely know that for whatever reason they are unable and/or unwilling to do so. So we're on our own.

Knowing all the things that are wrong – rampant spamming of code-identical extensions under different names, the return of previously removed extensions under different names, an essentially broken extension permissions system, totally bogus 5-star reviews, conscientious developer reports going completely unheeded, "Featured" extensions having no additional value whatsoever, and more – the title Wladimir gave to his extremely informative blog posting of *"Chrome Web Store is a mess"* seems entirely fitting.

I author these show notes in Google Docs every week. So I'm in a web browser while I'm writing this. And at one point while I was writing this, I looked up at the top of my browser with the intention to enumerate the browser extensions I'm using. Then I realized with a private smile that none of this applied to me, since I don't use Chrome at all – I'm happily using Firefox where the full strength uBlock Origin will continue to work.

While I'm sure that many of the same issues plague Mozilla's extension repository, Wladimir's comments did indicate that Mozilla and Opera have been much more responsive to abuse reports. And, if nothing else, it's Chrome that has by far the largest target painted on its back. In this case, I'd rather stick with the "also ran…" browser.