Security Now! #1061 - 01-20-26

More GhostPosting

This week on Security Now!

• RAM pricing to affect enterprise firewall equipment. • Anthropic provides sizeable support to Python Foundation. • The FTC clamps down on GM's secret sale of driving data. • "ANCHOR" replaces "CIPAC" for industry-government sharing. • Germany planning to legislate total access to global data. • Grubhub becomes the latest ShinyHunters extortion victim. • Let's Encrypt's 6-Day certs are available to everyone. • Iran plans to permanently take itself off the Internet. • HD Tune before and after a SpinRite Level 3 refresh. • Some great listener feedback, and • More trouble from GhostPoster malicious browser extensions.

"Never pay another DVD rewind fee again! Compatible with all disc formats: DVD-R, DVD-RW, DVD+R, DVD+RW, CDR, CDRW, Audio CD, VCD! Multi-region, code-free rewinder capable of rewinding all six region DVD's including RCE/REA encoded discs!"

"We have tested the DVD Rewinder with the next generation disc media including Blu-Ray®, and HD. The DVD Rewinder also works with Sony Playstation®, Xbox®, and other disc based console system media. The DVD Rewinder works with all disc based digital media to provide optimized digital experience. Visual indicators blink and audible sounds are played while your digital media is "reversed". The DVD Rewinder also has a USB port for MP3 players and USB media."

Security News

Enterprise firewall pricing expected to rise

Any of our listeners who provide purchase planning guidance for high-end network security products may wish to consider advising to have those purchases made sooner rather than later. Some recent commentary about the effect the rising cost of RAM will likely have on the security equipment sector suggested that prices could be expected to rise there as well. The commentary said:

> *The current price hikes and supply shortage of DRAM memory chips are expected to also impact firewall makers and the cybersecurity market. DRAM is a crucial component for the manufacturing of modern next-gen firewalls, a staple in the cybersecurity defense of any major enterprise.*
>
> *Investment advisory firm Wedbush says firewall companies will see thinner margins this year due to the rising DRAM costs. This will impact their bills of materials, with the extra costs being passed on to customers as product price increases. This will likely lead to lower sales, smaller profit margins, and weaker investor yields. Companies like Fortinet, Palo Alto Networks, and Check Point are expected to see the biggest headwinds on the stock market this year as a result of DRAM hikes.*
>
> *Firewall makers join laptop, PC, and smartphone vendors, all of which are expected to see big headwinds this year due to collapsing sales. DRAM prices have been up between 60% and 70% since last year and are expected to grow another 50% in the first quarter of the year alone.*
>
> *The production of most of this year's DRAM supply has already been purchased by AI companies for use in their future data centers. DRAM maker Micron has exited the consumer market and focused strictly on supplying AI and data center makers. South Korean company SK Hynix is also pondering a similar decision from both the DRAM and NAND/SSD markets.*

I've mentioned that I recently purchased my next small form factor desktop PC from Lenovo a few months before I plan to deploy it, due to the expectation that PC vendors will soon have no choice other than to raise their prices for full systems. And several months before that I had mentioned that I was similarly glad to have purchased replacement servers for GRC after the second of the five I had died. I wanted to be ready with replacement servers standing by in case I were to lose another. At the time, those server replacements were "just in case". But now I'm glad, since I always prefer to stuff my servers with as much RAM as they can handle, and last summer RAM was still amazingly inexpensive. But not so any longer.

So I think that the takeaway here, is that if someone already had plans to purchase high-end RAM-intensive network security equipment sometime soon, it might make sense to cut the purchase order soon, before prices rise as they are expected to.

The Python Software Foundation gains support from Anthropic

Last week, the Python Software foundation announced some very welcome financial support from Anthropic. Under the headline *"Anthropic invests $1.5 million in the Python Software Foundation and open source security"*, they wrote:

We are thrilled to announce that Anthropic has entered into a two-year partnership with the Python Software Foundation (PSF) to contribute a landmark total of $1.5 million to support the foundation's work, with an emphasis on Python ecosystem security.

This investment will enable the PSF to make crucial security advances to CPython and the Python Package Index (PyPI) benefiting all users, and it will also sustain the foundation's core work supporting the Python language, ecosystem, and global community.

Anthropic's funds will enable the PSF to make progress on our security roadmap, including work designed to protect millions of PyPI users from attempted supply-chain attacks. Planned projects include creating new tools for automated proactive review of all packages uploaded to PyPI, improving on the current process of reactive-only review. We intend to create a new dataset of known malware that will allow us to design these novel tools, relying on capability analysis. One of the advantages of this project is that we expect the outputs we develop to be transferable to all open source package repositories. As a result, this work has the potential to ultimately improve security across multiple open source ecosystems, starting with the Python ecosystem.

This work will build on PSF Security Developer in Residence Seth Larson's security roadmap with contributions from PyPI Safety and Security Engineer Mike Fiedler, both roles generously funded by Alpha-Omega. Anthropic's support will also go towards the PSF's core work, including the Developer in Residence program driving contributions to CPython, community support through grants and other programs, running core infrastructure such as PyPI, and more.

We could not be more grateful for Anthropic's remarkable support, and we hope you will join us in thanking them for their investment in the PSF and the Python community.

This is great and welcome news. $1.5 million likely makes a big difference to the Python project, as it would to any volunteer-driven open source effort. And, given the insane flows of cash the AI sector is seeing, $1.5 million doesn't even qualify as a drop in the bucket – it's more on the level of some vapor for the likes of any mainstream commercial AI vendor.

At the same time, much as this will be welcome support on the receiving end, we should also acknowledge that it's likely a clever investment on Anthropic's part. The line from the announcement that caught my eye was: *"Planned projects include creating new tools for automated proactive review of all packages uploaded to PyPI, improving on the current process of reactive-only review."*

"Automated Proactive Review." – in other words, deploying AI to examine all newly submitted Python Package code. And whose AI do you imagine the Python Software Foundation will choose to deploy? Given Claude's current code analysis strength, Authropic's solution would probably be the one to choose, even if they hadn't directly and deliberately aimed the PSF in that direction with their generous tax deductible gift. And if that work does turn out to be extensible to other open source ecosystems it could turn into a big win for everyone involved. So I say: "Good move, Anthropic!"

The FTC clamps down on GM's sale of their car owners' private driving data.
One of the more egregious privacy-invading behaviors that has come to light is the idea that car makers might be generating additional revenue for themselves, behind their car owners' backs, by selling data about their driving to insurance companies.

The question has been whether or not individual drivers may have consented to this. I would argue strongly that it's not possible to actually "consent" to something that's never explicitly described and explained, and which appears in a purchase agreement's legalize fine print. I've been driving for about 55 years, and I've purchased a few cars during that time. I've never attempted to read any of the fine print. I presume that as a U.S. consumer my rights will be protected by my government's agencies whose job it is to be a check on corporate greed, and to make sure that consumers who don't read the fine print get a fair shake nevertheless.

To that end, last Wednesday, the FTC posted an announcement under their headline: *"FTC Finalizes Order Settling Allegations that GM and OnStar Collected and Sold Geolocation Data Without Consumers' Informed Consent"* They wrote:

> *The Federal Trade Commission finalized an order with General Motors and OnStar settling allegations that they collected, used, and sold consumers' precise geolocation data and driving behavior data from millions of vehicles without adequately notifying consumers and obtaining their affirmative consent.*
>
> *Under the order finalized by the Commission, General Motors LLC, General Motors Holdings LLC, and OnStar, LLC (collectively GM), which are owned by General Motors Company, are prohibited from sharing certain consumer data with consumer reporting agencies. They also are required to take steps to provide greater transparency and choice to consumers over the collection, use, and disclosure of their connected vehicle data.*
>
> *In a complaint first announced in January 2025, the FTC alleged that GM used a misleading enrollment process to get consumers to sign up for its OnStar connected vehicle service and OnStar Smart Driver feature. The FTC also alleged that GM failed to clearly disclose that it collected consumers' precise geolocation and driving behavior data via the Smart Driver feature and sold it to third parties without consumers' consent.*
>
> *The final order approved by the Commission imposes a five-year ban on GM disclosing consumers' geolocation and driver behavior data to consumer reporting agencies. This fencing-in relief is appropriate given GM's egregious betrayal of consumers' trust. And for the entire 20-year life of the order, GM will be required to:*
>
> - *Obtain affirmative express consent from consumers prior to collecting, using, or sharing connected vehicle data (including sharing data with consumer reporting agencies), with some exceptions such as for providing location data to emergency first responders;*
> - *Create a way for all U.S. consumers to request a copy of their data and seek its deletion;*
> - *Give consumers the ability to disable the collection of precise geolocation data from their vehicles if their vehicle has the necessary technology; and*
> - *Provide a way for consumers to opt out of the collection of geolocation and driver behavior data, with some limited exceptions.*
>
> *The Commission voted 2-0 to approve the final order and complaint as well as provide responses to commenters.*

In addition to General Motors, Hyundai has been found to be sharing its drivers' data with Verisk, one of the major brokers for this information, and both Honda and Toyota are believed to be doing the same.

This nauseating spying on the part of automakers feels so similar to the idea of consumer ISPs surreptitiously monitoring and tracking their own subscriber's Internet usage and behavior

without their knowledge and permission. Again, the fact that some fine print may technically give them legal permission doesn't make it right.

Leo used to introduce me on this podcast as the person who coined the term "Spyware" and who created the world's first spyware removal tool. Both of those things are true. I named that first anti-spyware utility "OptOut" and I will never forget the raw fury that was expressed in the email end users were sending to the spyware parent's company, Aureate. It was way over the top, but that's how people reacted to the affirmative discovery of secretly installed spyware residing inside their machines. It was never my intention to put Aureate out of business. But it turned out that their entire business model was only viable while they remained unknown and secretive. Once people learned about them no one wanted anything to do with them.

My creation and publication of OptOut generated so much antipathy toward them that I spoke to their leadership on several occasions. I came to understand that they were not bad people. The Aureate system was a revenue-generation library that shareware and freeware authors could add to their software to display advertisements on the app's UI surface. So the Aureate system was supposed to "advertising enable" shareware to generate some revenue from the software's use.

The BIG mistake Aureate made was in relying upon the freeware and shareware authors to notify their users that this was taking place. No authors did that. Or if they did it was buried deep down in the software's license agreement that no one ever read or understood. I explained to the Aureate guys that they needed to take independent responsibility for their system by displaying their own permission dialog to get the end user's permission. Most of the anger – and oh boy was it palpable! – was over the fact that this was going on behind the users' backs.

So, the names have changed but the behavior has not. GM knows that if their users were clearly asked whether they would like to have detailed data about their driving habits sold for profit to third parties who will then resell it to their insurance provider to justify increases in their insurance rates, no one would say "You betcha! Sign me up for that!" Similarly, ISPs know that no one would want to have their detailed use of the Internet resold to data brokers.

We know that the opinions and votes of our politicians can be deeply influenced by commercial interests through lobbying. So we need to have independent consumer watchdog agencies such as the FTC watching our backs for us.

ANCHOR to replace CIPAC

Last year we touched upon the crucial need for industry executives to be able to disclose known security incidents to government officials without fear of reprisals from the government. This was the critical role that CIPAC had. CIPAC stood for the Critical Infrastructure Partnership Advisory Council. Last Wednesday, the publication CyberScoop published a very nice piece about the pending replacement agency. They wrote:

> *The Department of Homeland Security is finalizing plans for a new body that would replace the functions of the Critical Infrastructure Partnership Advisory Council (CIPAC) and serve as a communications hub between industry and government to discuss ongoing threats to U.S. critical infrastructure, including from cyber attacks.*
>
> *Under previous administrations, CIPAC served as a nerve center for federal agencies, industry and other stakeholders. While industry widely praised its utility, the council was one of many DHS advisory bodies that were shuttered last year by Secretary of Homeland Security Kristi*

Noem after President Donald Trump returned to office.

Now, according to multiple sources, a proposed regulation for a new replacement council is in the final stages of review and approval from Noem's office.

The new body will be called the Alliance of National Councils for Homeland Operational Resilience, or "ANCHOR," and will also serve as an umbrella organization for other federal sector risk management agencies. Its goal is to restart conversations and planning around infrastructure security that took place under the previous CIPAC, according to a former DHS official.

The official, who requested anonymity to discuss the administration's plans, said all 15 federal sector coordinating councils have been briefed on ANCHOR. One of the primary differences between CIPAC and ANCHOR will be in structural authorities and liability protections.

CIPAC was essentially "an advisory council that could be chartered to create other advisory councils" that needed Secretary-level approval and contained rigid rules requiring separate charters for every new council that was then stood up.

This created "a waterfall effect" of bureaucracy that made CIPAC a poor vehicle for holding broad conversations between not just DHS and industry, but all other federal sector risk management agencies and sector coordinating councils.

The official said: "What DHS has strived to do is create a new framework for engaging on threat conversations and pre-deliberative policy conversations impacting security outcomes with sectors and the private sector, without having to create all these waterfall advisory councils or new charters and all that stuff."

Okay. So far so good. Any reduction in needless bureaucracy sounds like a good thing. CyberScoop's reporting continues:

Under CIPAC, conversations between government and industry were also "closed by default" to the public, with mandatory liability protections for every conversation and setting. Often, the most the government could do was issue a press release or cite comments under Chatham House Rule. Under ANCHOR, there is expected to be wider latitude for DHS or other councils to open certain meetings to the public, or provide transcripts of conversations they hold with stakeholders.

However, the official emphasized that liability protections remain one of the last unresolved issues. The administration is still determining when those protections would or would not apply to ANCHOR-related discussions between government and industry and further changes could be made to assuage industry.

Other federal laws, such as the Cybersecurity and Information Sharing Act of 2015, only provide liability coverage for "one to one" conversations between a company and the government. The previous entity, CIPAC, by contrast, provided a liability shield for "one-to-many" engagements, where a company may engage with federal, state and local agencies as well as other companies and entities.

The official said: "That created a well understood and important liability shield which allowed senior officials, all the way up to the CEO of private sector companies, to openly communicate

with each other."

Following the initial publication of this reporting, a DHS spokesperson in a statement did not dispute the description of ANCHOR provided by CyberScoop, but called discussions of an imminent regulation release "premature." The spokesperson said: "We look forward to sharing more details once we have something to announce."

This week, Adrienne Lotto of the American Public Power Association told Congress that liability protections in CIPAC were critical to fostering open dialogue between industry and government around cybersecurity and infrastructure protection. She also signaled that a new advisory council was forthcoming, saying industry "was apprised by DHS that the administration's proposed CIPAC replacement is ready for publication in the Federal Register" while encouraging the administration to finalize the plans "quickly."

Even with some uncertainty around ANCHOR's structure and liability protections, many industry executives are likely to embrace the return of information-sharing partnerships that they believe were vital to understanding the digital and physical threat landscape facing their industry sectors. Last year, industry groups lamented the disbanding of CIPAC to members of Congress, prompting Rep. Andrew Garbarino, now chair of the Homeland Security Committee, to pledge he would "look into this and hopefully speak to the administration to try to fix this."

The former DHS official said they expected ANCHOR to be largely welcomed by many industries who have called for the restoration of CIPAC, even as they look to grapple with the Trump administration's new approach. The official said: "Everybody who wants to talk in groups is going to be excited to have it back. At the same time, those who are concerned about the amount of risk it opens up will want to see the details."

I clearly recall being concerned by the disbanding of the original CIPAC, since there are clearly things the government alone can do, which private industry may need. But if a fear of the consequences of divulging serious problems keeps industries silent, that's not good. I like the sound of an improved structure that sidesteps the need to design and spawn endless subcommittees. And it sounds as though the need for liability protections is clear. So let's hope we soon get ANCHOR up and running.

Germany, WTF?

I have some reporting that was obtained from translations from German. At this point today, since it describes Germany's new legislation as "pending" as opposed to "enacted", I didn't want to spend time digging into source material – which would all need to be translated. And also, my assumption is that if or when this does occur it will have plenty of multi-sourced coverage.

So today, I'm just going to share the reporting that I have, and everyone will quickly see why it was worth sharing it "as is" for now. The reporting read:

German lawmakers are working on a new law that will grant the country's intelligence agency new and extensive hacking and surveillance powers. The primary intent of the new law is to free up the Bundesnachrichtendienst (BND) from relying on the US National Security Agency (NSA) for threat information and bring Germany's interception capabilities on par with other European countries, such as France, Italy, the Netherlands, and the UK.

According to a draft of the new law obtained by German media, the BND will have the power to intercept full internet communications and not just metadata as it is allowed today. The agency will also be allowed to store the data for up to six months, which will allow it to better index and search it for threat intel. The BND will also have its offensive hacking mandate extended. The law will allow the agency to hack foreign internet service providers and retrieve information about its targets if the companies do not cooperate or provide the requested data.

According to reports, this provision will apply to major US companies and infrastructure providers like Google, Twitter, and Meta, which have been known to be prickly around surrendering such information in the past. The agency could previously intercept the communications of individuals abroad, but now the BND will also be allowed to put any foreigner in Germany under surveillance. The same goes for journalists working for foreign state-run media organizations, which German lawmakers say are acting more like "agents" of a foreign state than independent reporters.

Finally, BND agents will also be allowed to enter apartments and deploy their "federal trojan" on a target's device. According to reports, the new law's draft is 139-pages long and almost doubles the BND's previous capabilities.

I think the short version of what this means is “thank goodness for state-of-the-art encryption” which we have every reason to believe is utterly unbreakable by anyone.

And while Germany’s legislation might at first seem like egregious overreach, we know that the U.S.’s National Security Agency has already built a massive data center of over 1 million square feet about 20 miles South of Salt Lake City, Utah. And while the details are kept close, it’s well known to be a massive data **storage** facility. We’ve often noted that there may be value in storing massive quantities of encrypted data that cannot be deciphered today, but may be decipherable using tomorrow’s technology. So it’s easy to imagine that the internal encrypted communications of the U.S.’s global adversaries may be tapped and tagged and sent to Utah for long term archiving. And then, once the NSA’s quantum computing technologies come online in the future, the public key crypto handshakes that established the ephemeral secret symmetric keys might be broken and those communications, even though by then no longer current, might be important to obtain.

I sometimes feel that the EFF’s, the Electronic Freedom Foundation’s, absolutism about privacy rights and encryption goes overboard. But when we see examples like this, of how aggressively foreign governments and our own are pursuing information that for the most part they have no true need for, I appreciate that the EFF is working to always provide some counter-pressure.

Someone inside Grubhub clicked the wrong link.
We appreciate that it can happen to anyone. And now it appears that someone inside Grubhub clicked a link they should not have, which permitted the infamous “ShinyHunters” gang to obtain authentication credentials. BleepingComputer, which reported on this exclusively last Thursday, headlined their reporting: *“Grubhub confirms hackers stole data in recent security breach”*. They wrote:

Food delivery platform Grubhub has confirmed a recent data breach after hackers accessed its systems, with sources telling BleepingComputer the company is now facing extortion demands.

Grubhub told BleepingComputer: "We're aware of unauthorized individuals who recently downloaded data from certain Grubhub systems. We quickly investigated, stopped the activity, and are taking steps to further increase our security posture. Sensitive information, such as financial information or order history, was not affected."

Grubhub would not respond to any further questions regarding the breach, including when it occurred, whether customer data was involved, or if they were being extorted. However, the company confirmed that it is working with a third-party cybersecurity firm and has notified law enforcement.

Last month, Grubhub was also linked to a wave of scam emails sent from its b.grubhub.com subdomain that promoted a cryptocurrency scam promising a tenfold return on Bitcoin payments. Grubhub said at the time that it contained the issue and took steps to prevent further unauthorized messages, but would not answer further questions related to the incident. It is unclear whether the two incidents are connected.

*While Grubhub would not share further details, multiple sources have told BleepingComputer that the ShinyHunters cybercrime group **is** extorting the company. BleepingComputer attempted to verify these claims with the threat actors, but they, too, refused to comment.*

I'll interject that the threat actors' silence at this juncture is to be expected, since part of their promise in return for receiving an extortion payment will be their silence. Since they presumably still hope that the returns from their data breach will result in a payday, much as they have shown a willingness to brag, they're certainly not going to talk to the press until it's clear that doing so would not compromise their negotiations and their extortion payment, if any. BleepingComputer continues:

According to sources, the threat actors are demanding a Bitcoin payment to prevent the release of older Salesforce data from a February 2025 breach and newer Zendesk data that was stolen in the recent breach.

That all tracks the reporting we've been doing here, where we noted a month or two ago that the ShinyHunters gang had switched to attaching ZenDesk users after they had apparently fully played out their multiple earlier Salesforce breaches. BleepingComputer wrote:

Grubhub uses Zendesk to power its online support chat system, which provides support for orders, account issues, and billing. While it is unclear when the breach occurred, Bleeping-Computer was told that it was through secrets & credentials stolen in the recent Salesloft Drift data theft attacks.

In August, threat actors used stolen OAuth tokens for Salesloft's Salesforce integration to conduct a data theft campaign between August 8 and August 18, 2025. According to a report by Google's Threat Intelligence team (Mandiant), the stolen data was then used to harvest credentials and secrets to conduct follow-up attacks on other platforms. Google reported "GTIG observed UNC6395 targeting sensitive credentials such as Amazon Web Services (AWS) access keys (AKIA), passwords, and Snowflake-related access tokens." ShinyHunters claimed at the time to be behind the breach, stating they stole approximately 1.5 billion data records from the "Account", "Contact", "Case", "Opportunity", and "User" Salesforce object tables for 760 companies.

Let's Encrypt's 6-Day Certs are Online
Last Thursday, January 15th, Let's Encrypt announced under the headline: *"6-day and IP Address Certificates are Generally Available"*. They wrote:

> *Short-lived and IP address certificates are now generally available from Let's Encrypt. These certificates are valid for 160 hours, just over six days. In order to get a short-lived certificate subscribers simply need to select the 'shortlived' certificate profile in their ACME client.*
>
> *Short-lived certificates improve security by requiring more frequent validation and reducing reliance on unreliable revocation mechanisms. If a certificate's private key is exposed or compromised, revocation has historically been the way to mitigate damage prior to the certificate's expiration. Unfortunately, revocation is an unreliable system so many relying parties continue to be vulnerable until the certificate expires, a period as long as 90 days. With short-lived certificates that vulnerability window is greatly reduced.*
>
> *Short-lived certificates are opt-in and we have no plan to make them the default at this time. Subscribers that have fully automated their renewal process should be able to switch to short-lived certificates easily if they wish, but we understand that not everyone is in that position and generally comfortable with this significantly shorter lifetime. We hope that over time everyone moves to automated solutions and we can demonstrate that short-lived certificates work well.*
>
> *Our default certificate lifetimes will be going from 90 days down to 45 days over the next few years, as previously announced.*
>
> *IP address certificates allow server operators to authenticate TLS connections to IP addresses rather than domain names. Let's Encrypt supports both IPv4 and IPv6. IP address certificates must be short-lived certificates, a decision we made because IP addresses are more transient than domain names, so validating more frequently is important. You can learn more about our IP address certificates and the use cases for them from our post announcing our first IP Certificate.*
>
> *We'd like to thank the Open Technology Fund and Sovereign Tech Agency, along with our Sponsors and Donors, for supporting the development of this work.*

As I've said before, the shortening of the maximum lifetime of web server DV domain validation certificates will eventually drive GRC to use Let's Encrypt's free certificates. Once I switch to their solutions I will definitely establish a periodic voluntary payment to them, much as I have with Wikipedia, since I feel that it's important to support the infrastructure that makes that possible, even if the entire necessity of any of this is something I could not disagree with more. So be it.

It's never been clear to me who has such a huge problem holding onto their web server's private keys. All indications are that the entire thing is a made-up problem. Remember that even if a bad guy could somehow arrange to obtain a valuable domain certificate, it's not as if just having that in any way allows them to impersonate the target site. They must still somehow arrange to cause their victim's Internet traffic to believe that it's going to the real domain's IP address, while it is, instead, re-routed to a spoofed server. So either a DNS compromise must also be effected, or physical interception and re-routing must be achieved.

If this was EVER happening – if it ever happened – it would be huge news. We would know about it. Instead, it's crickets.

And I get it that the Let's Encrypt guys need to say that revocation is broken, but it no longer is.

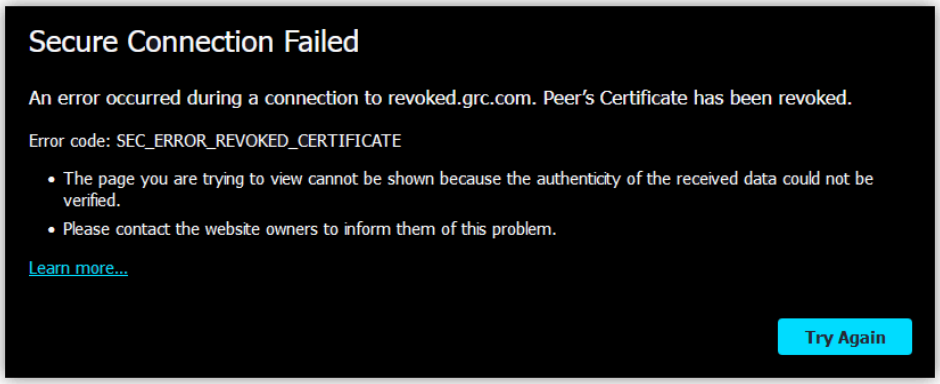

Any of our long-term listeners know that I was onto all of this, pointing this out and drawing attention to this as loudly as I could, before anyone else was doing so. I looked foolish at the time, like I was tilting at windmills – what's the big deal, Gibson? I created "revoked.grc.com" to clearly demonstrate that none of this was working. It is now, everywhere. And it's even been solved on the client thanks to Bloom filters.

And just that I'm clear, I think it's truly great that Let's Encrypt is not offering 6-day TLS DV and IP validated certificates – for those who feel that they need them. That's great. It's the being forced to use shorter life certificates – whether for the web or for code signing – that feels so wrong and regressive to me. I don't need a nanny. Few of us do. And as I've said, if anyone did it would be making news.

Iran plans to permanently disconnect its citizens from the Internet

Several news outlets reported on something that caught my attention, mostly because it's so sad and, in my own opinion, so wrong-minded. The news is that Iran plans to extend its current disconnection from the Internet, which began on January 8th, to make it permanent. Technical reports have indicated that efforts are being made to restrict the use of messaging apps for internal use only. All satellite dish antennas are being gathered up and technology is being finalized to identify network traffic that transits across Starlink and other space-based providers.

Iran's ruling theocracy is what it is. It's been clear that the influence of the West, largely though obviously not exclusively brought to Iran by the Internet, has been a challenge to the nature of its historical theocratic rule. But Iran's population today is not old. Its median age is somewhere between 33 to 34. So half of Iran's population is younger than 33 to 34 and currently about one quarter of Iran's population are children under the age of 15. Cutting that population off from all external Internet access certainly seems destined to fail.

But what appears not to be destined to fail is any mass storage drives that are given to SpinRite:

SpinRite

Donn Edwards

Dear Steve: You have often mentioned how SpinRite improves SSD performance, and we have seen the results of its benchmark tests. But here is a different view. My friend panicked when his computer wouldn't boot. It has a Crucial 480GB SSD boot disk, and a Seagate 1TB HDD data disk. Not knowing whether the problem was hardware related or not, I rescued the drives and connected the SSD to my own desktop PC to see if the data was intact. All appeared fine, so I ran HD Tune to look at the SMART data and run its benchmark.

The drop in performance (shown in the HD Tune Pro chart) , particularly at the start of the drive, was troubling, so I ran SpinRite 6.1 on level 3 and it took around 3 hours. I could see it having trouble writing to the drive, but in the end no data was lost.

Afterward (see the post-SpinRite chart on right) it's clearly fixed. I backed up all the data files from his HDD and put both drives back in the PC. When we plugged in all the cables and screens, his PC worked! So whether it was the SSD, or a bad cable connection, or something else, I don't know. But what I do know for sure is that his SSD is working much better than before. The graphs show it, and he is very relieved.

Keep up the good work! / Donn Edwards / Johannesburg, South Africa

Listener Feedback

Jeff Ekstrand

You can find the advertising ID on Roku via some secret menus. On the remote you can do some convoluted button pushes to access these menus. One of them contains the advertising ID (I do not remember which one).

Cheat Sheet!

#1 Developer settings: Press: Home × 3 - Up × 2 - Right - Left - Right - Left - Right
#2 Wireless secret screen: Press: Home × 5 - Up - Down - Up - Down - Up
#3 Secret screen: Press: Home × 5 - Fast Forward × 3 - Rewind × 2
#4 Secret screen 2: Press: Home × 5 - Up - Right - Down - Left - Up
#5 HDMI secret screen: Press: Home × 5 - Down - Left - Up × 3
#6 Platform secret screen: Press: Home × 5 - FF - Down - Rewind - Down - FF
#7 Channel info menu: Press: Home × 3 - Up × 2 - Left - Right - Left - Right - Left
#8 Reboot shortcut: Press: Home × 5 - Up - Rewind × 2 - Fast Forward × 2

The link below is the YouTube video that has the information: https://youtu.be/vxDOu7unNX8

You can also search for "Tech It Before You Wreck It" channel and video titled "The 8 ROKU Secret codes they DON'T want you to know about". Thanks for the work that you put into the podcast. It is my favorite podcast. /Jeff Ekstrand

It's VERY cool, it works, and any techie listener of ours with a Roku will probably get a kick out of it. I found my Roku's clearly labeled Advertising ID on Secret Screen #2, which is displayed by pressing the Home button 5 times in a row followed by Up, Right, Down, Left, Up. The last item on that screen was my Advertising ID as a GUID-format identifier. There is a bunch of other information. As is generally the case with these sorts of hidden Easter Eggs, the information revealed probably won't change your life, but it does expose a bunch of the Roku's internals. The developer page allows you to enable the Roku's normally disabled internal web server which you can then access with any web browser on your LAN.

The Cheat Sheet Jeff provided in his email is present in the notes for the 11-minute YouTube video which has had more than 1.2 million views since being posted two months ago on November 19th.

Michael Wright

Hello Steve, I'm a first-time emailer to you who has been listening to your show for a couple of years now and find it a great resource for keeping up with developments in the world of cybersecurity. Thanks so much for doing these podcasts!!

I'm a week behind with the podcasts and today finished last week's podcast. You made a good point about how there should be no legitimate reason for anyone to have their MongoDB server accessible over the Internet. That got me wondering if people are deploying MongoDB servers without even realizing they are publicly accessible. I'm referring to cloud deployments, where for many flavours of deployment, a public IP address is automatically created. With traditional On-Prem, making a server accessible over the internet required work to be done. For example creating a NAT rule on a firewall to translate a public IP address to a private IP address, however, with public cloud this is often done automatically. If people are deploying systems to the cloud without having an understanding of cloud deployment and how this differs from on-prem, I could certainly see how it could be possible to deploy a system without realising you just made it accessible to anyone anywhere on the internet. Would be interesting to know how many of the 86,000 exposed servers are using IP addresses reserved for public cloud.

Keep up the great work! p.s. On the topic of British time travel series, I found 'Bodies' to be a pretty good effort. Certainly a different take on the subject. Not sure if you have seen that one. /Regards, Michael

I suspect Michael is right, and that many of these MongoDB server instances are spun-up in the cloud. Although this may be an explanation, it certainly isn't any excuse. What's happening is very wrong. So the question is, how did this happen? It's likely a case of the user assuming that "those in charge" are doing the right thing, whereas "those in charge" wrongly assume that their users are aware of the implications of spinning up random server instances in the cloud and they assume that those users will prevent public exposure if they don't want it.

In other words, one hand doesn't know what the other one is doing and they each assume that the other one is taking responsibility for the expected (and needed) network security. The problem is that those who design these cloud-based services heavily promote their super-easy "one-click" server activation. So they're offering their inherently insecure solutions to a level of user who has very little comprehension of the full implications of clicking that "please create a MongoDB server instance for me" button.

I wanted to focus upon this specific instance because I suspect that this lack of communication with its assumption that the other party will take care of securing things has long been a major source of network insecurity. Several months ago I noted that the early Cisco routers, which had no built-in notion of public-facing WAN interfaces versus private-network LAN interfaces, treated all of their network interfaces identically. Those early routers also had their various network services enabled out of the box. Back then, for example, you had to manually add a "nohttp" command to the router's start-up configuration script if you did not want the router's built-in http server to be running by default. I very clearly recall needing to deliberately turn off a handful of services that I had no need for and didn't want to have running every time to set up a new Cisco router. The engineer/designers of these early routers must have assumed that their devices would only and always be used by other expert network engineers. And since Cisco was always selling the "security" of their products, non-expert purchasers reasonably assumed that Cisco would have their back and that the router's operation would be secure out of the box, when it was anything but. Instead, it was bristling with enabled and insecure gee-whiz features that were entirely peripheral to the router's operation. The lesson here is that each side's assumptions about the other were wildly incorrect and serious vulnerabilities resulted.

Bob's note's subject was: *"**Cyber Attack** - Was My Experience Unique?"*

Hello GRC Team, I have been a big fan and Spinrite customer since learning about your Spinrite product on The Tech Guy.

Recently, I experienced a type of cyber attack I had not heard of. I can go into more detail, but basically, a software program, Screen Connect, was remotely installed on my PC and launched with no interaction by the client (me). I became aware of the attack when I was at my mom's house and my phone started notifying me of money transfers that I did not initiate. I freaked out as you might imagine.

I rushed home and when I got there, I found that my machine had been hijacked. My screens were blacked out with "Screen Connect" in large white letters. I was unable to do anything other than shutting down the machine.

Needless to say, I have been dealing with the aftermath and fortunately I am not out too much money, but I found out who my friends and foes are in terms of how they did or did not help me cancel the transactions. In short, PayPal's response was abominable.

I assume the criminal used a sniffer to find my IP address and since my machine was idle,

they were able to install and launch Screen Connect without detection (no client interaction to install and launch the software is considered a feature of the product). IMO, the software is like a gun, misuse can lead to devastating results! They offer a free 15 day trial, but I didn't check to see if it is full featured.

What do you think about this? Short of keeping my machine powered off, what could I have done to block this type of attack? Any insight would be appreciated! Regards, Bob.

Yikes! This is the nightmare scenario for any individual. I've omitted Bob's last name to protect his identity. No one wants to be required to authenticate with every service we use, every time we use them. So being persistently logged into many services is the choice most of us make. But with that convenience comes the consequence that anyone – and anything – that's able to use our persistently logged in computer can act on our behalf.

The abuse of persistent logon is what bit Bob. Bob doesn't know, so we don't know, exactly how someone managed to crawl into his PC. Through the years of this podcast we've seen many different ways this could have happened. But by far the most likely is that Bob or someone using Bob's computer clicked on a malicious link. Last week, Leo shared his own incident which forced him to cancel and have two credit cards reissued. And I mentioned that I had received a text message that I briefly considered to be valid because by pure chance it fit into the context of my life and it made sense to me.

So it's certainly not the least bit far-fetched to imagine that Bob, or someone who uses Bob's PC, might have made the mistake of clicking on a malicious link in email. That's all that's needed. That could have established an outgoing connection to an attacker who was then able to install the "client-free" Screen Connect remote control software. The attacker could then have waited until the PC that had been left running and unattended, showed no use of its keyboard or mouse for some period of time, then took the opportunity to begin sending the owner's money to remote accounts. For example, PayPal allows zero-authentication transfers of cash from the bank accounts and credit cards associated with the user's PayPal account.

When Leo and I speak to the attendees of ThreatLocker's Zero Trust World conference in Florida this coming Wednesday, March 4th, our discussion will be titled "The Call is Coming from Inside the House." We're going to be talking about the growing need for enterprises to actively protect themselves from anything their own employees might do. Whether it's deliberate or inadvertent doesn't matter since the result to the enterprise is the same either way.

Doing this effectively means imposing significant limitations upon everyone who has access to the enterprise's internal network. I'll be arguing that while it will not be at all easy, there is no longer any other way to further increase security. Given everything we've seen in the past year, it's clear that the spoofing of enterprise employees is the next "big growth" threat vector.

But for the individual PC user at home, no one wants to impose severe restrictions upon themselves when they're working within their own safe enclave at home. I certainly wouldn't. In this case, this happened to Bob because his PC was able to act without his physical presence to send his money away. The practical solution to this would be the inclusion of a simple biometric authentication for anything that requires Bob's presence. Having a fingerprint reader integrated into our keyboards or mice to confirm the identity of anyone who is requesting a protected action would prevent these sorts of unattended or other-attended attacks.

And, for example, the Bitwarden password manager – a sponsor of the TWiT network – fully supports unlocking with biometric authentication on Windows, macOS and Linux, and with all

Chromium based browsers, Firefox and Safari. So setting this up would certainly be possible.

But this still leaves user-spoofing as a problem. The most reasonable explanation for how the Screen Connect remote control software found its way onto Bob's machine is that he took some deliberate action – downloaded and installed something that incorporated this malicious functionality – without realizing it. So even biometric authentication would not have prevented that.

Rob Sherman – Subject: Feedback on Claude AI

Hi Steve, I just finished listening to last week's SN episode and as someone who has been using it constantly since the update came out I wanted to give you some feedback.

In short: It is absolutely insane how good it is. I am a Product Manager and not a programmer. So when my CTO told me that I needed to try it I wasn't sure why. I am now.

I had an internal project that I have been waiting to get programmer resources for over 6 months. Once I got Visual Studio setup with Copilot I gave it my product brief and after answering a few simple questions that Claude had, it began coding. An hour later I had a fully functioning Alpha. It did all the coding, designed and built a UI, and implemented a scanner to get all the data out.

Since then, when I have a few hours I'll just go in and tweak it. That dark mode I have been asking for the last year, its in there (The toggle for it is labeled "I finally got my dark mode) The build reporting and error checking I was told we wouldn't be able to do. It's done.

I have also completed 3 other projects that we weren't supposed to get to until Q3. It is amazing. I am so sold on it that I got myself a personal license and this weekend did a writeup on the E-drum application I have been waiting for someone to build. I gave it to Claude and now have my very own Alpha version.

That is not to say that it has been 100% smooth sailing. There is a learning curve to Claude especially, and I have blown through 200% of my monthly request at work in 14 days.

A few tips for anyone looking to get started with this:

1. Your individual "chats" with Claude have a size limit. Once you hit that limit you have to start a new chat. If you are just asking it a simple question you will be fine but any larger projects you will run out of room. I recommend starting any project by having Claude writeup a programming plan and tracking document. Then have it keep those files updated. That way if you have to start a new chat you can tell it to go read those docs to get up to speed.

2. Claude (In Visual Studio Copilot) won't let you upload PDF or other docs, but you CAN add MD files. I have taken to having Chat GPT summarize any files and turn them into MD format which I can then put in my project repo. Once in there, Claude is all set.

*3. Claude WILL lie to you. It is **always** a good idea to have it double check its own work. I had it write a bunch of new code. When it was done I told it "Hey would you take a look at this new code and check it for errors?" It found 4 items that needed fixing :)*

Thanks for everything you do / Rob / PS: Started taking Magnesium last week

More GhostPosting

GhostPoster lives on

Our final podcast of 2025 was titled "GhostPoster." For the short summary at the top of the show notes I summed it up, writing: "How a PNG Icon was used to infect 50,000 Firefox users." The discoverer of 17 different malicious Firefox add-ons was Koi Security. They discovered that PNG icon files were being used to contain and infiltrate obscured JavaScript into user PCs. Some of the extensions were FreeVPNs and others were junk extensions that someone who just wanted to collect free browser add-ons might add to their browsers. Nevertheless, more than 50,000 Firefox users had this malicious code running inside their browsers. So one of our takeaways was to avoid collecting crap from obscure sources that's not really needed. And, by the way, the phrase "FreeVPN" is an oxymoron.

So that was episode #1057. Why are we back here four weeks later for episode #1061? It's because following Koi Security's discovery, a different firm, LayerX, has reported their discovery of an additional 17 of the same, but this time they're not just attacking Firefox. Users of Edge and Chrome turn out to have been even earlier targets, with more than 840,000 downloads and installations.

LayerX's disclosure headline was *"Browser Extensions Gone Rogue: The Full Scope of the GhostPoster Campaign"* Here's what we now learn from LayerX's follow-on research. They write:

> Last month, researchers at Koi Security published a detailed analysis of a malicious Firefox extension they dubbed GhostPoster – a browser-based malware leveraging an uncommon and stealthy payload delivery method: steganography within a PNG icon file. This innovative approach allowed the malware to evade traditional extension security reviews and static analysis tools.
>
> Following their publication, our investigation identified 17 additional extensions associated with the same infrastructure and tactics, techniques, and procedures (TTPs). Collectively, these extensions were downloaded over 840,000 times, with some remaining active in the wild for up to five years. The GhostPoster malware employs a multi-stage infection chain designed for stealth and persistence:
>
> - Payload Encoding: The initial loader is embedded within the binary data of an extension's PNG icon.
> - Runtime Extraction: Upon installation, the extension parses the icon to extract the hidden data, a behavior that deviates from typical extension logic.
> - Delayed Activation: The malware delays execution by 48 hours or more, and only initiates C2 communication under specific conditions.
> - Payload Retrieval: The extracted loader contacts a remote C2 server to download additional JavaScript-based payloads.
>
> After activation, the malware is capable of:
>
> - Stripping and injecting HTTP headers to weaken web security policies (e.g., CSP, HSTS).
> - Hijacking affiliate traffic for monetization.
> - Injecting iframes and scripts for click fraud and user tracking.
> - Programmatic CAPTCHA solving and injection of additional malicious scripts for extended control.

These features indicate the campaign is not only financially motivated but also technically mature, emphasizing operational stealth and longevity.

The infrastructure uncovered by Koi Security was linked to 17 Firefox extensions, all sharing similar obfuscation patterns, command & control behavior, and delayed execution strategies. Our automated Extension malware lab feature confirmed the same threat actor infrastructure was also used to distribute extensions on the Google Chrome and Microsoft Edge Add-ons Store. Our analysis shows the campaign originated on the Microsoft Edge browser, with later expansion to Firefox and Chrome.

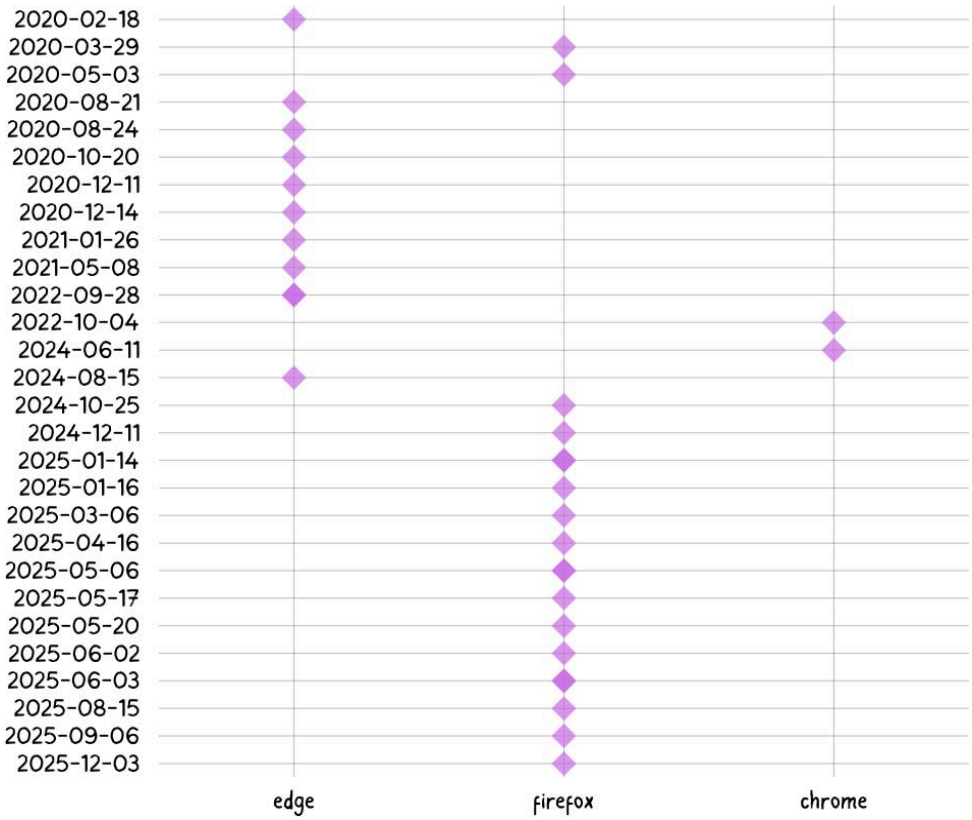

They provide a timeline chart which shows that the first known extension infected Edge browser users back in February of 2020. About six weeks later, at the end of March, Firefox was first hit. It was hit again at the beginning of May. Then a run of eight more malicious Edge extensions were released over the course of two years, from the end of August 2020 through the end of September 2022. Then, a month later, at the start of October 2022 the first Chrome extension was created. Then things were quiet for nearly two years until another edge extension appeared in August of 2024. But then after that it was all Firefox from the end of October of 2024 to today.

It's interesting that throughout all this time, only two known malicious extensions were seen to affect Chrome. It would be interesting to know why, since that's clearly the largest potential source of user installations. But in any event, the LayerX people expanded upon Koi's earlier findings by reporting:

- 17 additional confirmed extensions, with infrastructure overlap and common loader patterns.
- More than an additional 840,000 cumulative installs across Firefox, Chrome & Edge browsers.
- Malicious presence dating back to 2020, indicating long-term operational success, bypassing

all major browsers stores' security checks .

- And malware variants using alternate delivery mechanisms which suggests ongoing experimentation and adaptation.

Beyond the previously identified extensions, we observed a more sophisticated and evasive variant associated with the same campaign, which alone accounted for 3,822 installs:

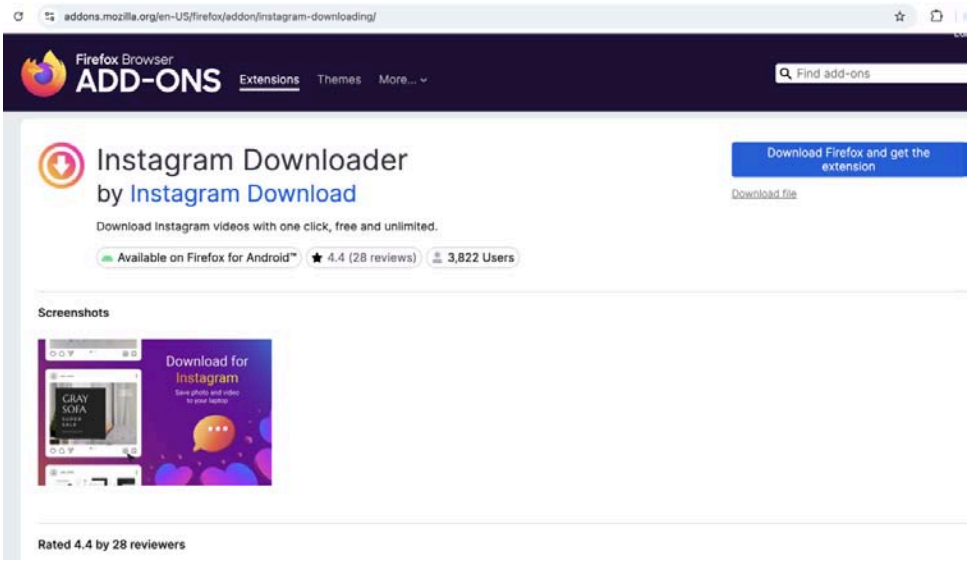

In this iteration, the malicious logic is embedded within the background script and leverages an image file bundled inside the extension as a covert payload container. At runtime, the background script fetches the image and scans its raw byte sequence for the delimiter [62,62,62,62] – corresponding to the ASCII string '>>>>'. All data following this marker is decoded as text and stored persistently in chrome.storage.local under the key instlogo. The stored data is later retrieved, Base64-decoded, and dynamically executed as an additional JavaScript payload.

This secondary script introduces further evasion by sleeping for approximately five days before initiating network activity. Upon activation, it fetches content from a remote server, extracts server-supplied data stored as Base64-encoded keys, and executes the decoded content, enabling ongoing payload updates and extended control. This staged execution flow demonstrates a clear evolution toward longer dormancy, modularity, and resilience against both static and behavioral detection mechanisms.

While Mozilla and Microsoft have removed the known malicious extensions from their respective stores, extensions already installed on user systems remain active unless explicitly removed by the user. This persistence underscores the limitations of store takedowns as a containment strategy, particularly for malware employing delayed activation and modular payload delivery.

The "Page Screenshot Clipper" had only 86 downloads. The "Full Page Screenshot" had 2,000. "Convery Everything" (whatever that is) had 17,171. But the "Translate Selected Text with Google" had 159,645 downloads. The other names, just to give everyone a feel for what these nasties are, were: Youtube Download, RSS Feed, Ads Block Ultimate, AdBlocker, Color Enhancer, Floating Player – Picture-in-Picture Mode, One Key Translate, Cool Cursor, Google Translate in Right Click. By the way, that one was the #1 download of all time, seeing more than half a million downloads alone: 522,398. There was also Translate Selected Text with Right Click, Amazon Price History, Save Image to Pinterest on Right Click, and that Instagram Downloader

that had 3,822 downloads.

While we all might determine that something's fishy about an offer for a FreeVPN, the screen shot of the Instagram Downloaded looks entirely legitimate and I can imagine downloading it without ever being the wiser. So this is truly a problem.

One thing that puzzles me is LayerX's suggestion that the removal of extensions from the web store leaves any already-downloaded and installed extensions in place and dangerous. We know that all of the browser vendors have the ability to remotely disable any browser extensions that are found to be malicious. I suppose it might be the case that a malicious extension that its malicious publisher withdraws from the store might slip under the radar since it's no longer being offered.

It might be that the post-installation mechanisms theSe extensions use, by moving their later downloaded code into the browser's permanent store affords them some post-removal protection.

But the convincing appearance of that Instagram Download extension is somewhat unnerving. It's important to note that Koi was aware of around 50,000 downloads and installs because, for whatever reason, they apparently were not looking back far enough.

One of the important takeaways here is that we must always remember that we can never know what we don't know. There's no point in getting overly worked up over things we cannot control – nor excessively worrying over what we don't know, since we don't know what that might be. But after seeing that there are things that we once didn't know that we do now know, it makes sense to wonder what else we might not know and, if nothing else, to behave with an appreciation that there are certainly some things that remain unknown.

Unfortunately, what they may be, we don't know.

What we do know is that we'll be right back here next week – same time, same channel – to continue sharing – to the best of our ability – events of interest that have recently transpired, don't ya know!?