

# Security Now! #1057 - 12-23-25

## GhostPoster

### This week on Security Now!

- North Korea's profitable fixation on cryptocurrency.
- Amazon uncovers a cryptomining sneaking into customer clouds.
- Insecure Docker API servers are also hosting cryptominers.
- A new and truly massive SmartTV-based botnet discovery.
- DNS Benchmark's 4th release.
- Who, besides Let's Encrypt, offers free automated certs.
- Some interesting listener feedback.
- And how a PNG Icon was used to infect 50,000 Firefox users.

Rather than discarding the heat from a power-sucking Bitcoin mining rig... why not use it to heat your home?



# Security News

## North Korean Hackers go for the Crypto

The blockchain analytics company "Chainalysis" posted an interesting end-of-the-year piece last Thursday titled: "North Korea Drives Record \$2 Billion Crypto Theft Year, Pushing All-Time Total to \$6.75 Billion". Their article was lengthy, so I'm not going to share it all, but it provides a very interesting look into today's digital asset industry. They start with five bullet points:

- *North Korean hackers stole \$2.02 billion in cryptocurrency in 2025, a 51% year-over-year increase, pushing their all-time total to \$6.75 billion despite fewer attacks.*
- *The DPRK is achieving larger thefts with fewer incidents, often by embedding IT workers inside crypto services or using sophisticated impersonation tactics targeting executives.*
- *The DPRK shows clear preferences for Chinese-language money laundering services, bridge services, and mixing protocols, with a 45-day laundering cycle following major thefts.*
- *Individual wallet compromises surged to 158,000 incidents affecting 80,000 unique victims in 2025, though total value stolen (\$713M) decreased from 2024.*
- *Despite increased Total Value Locked in DeFi, hack losses remained suppressed in 2024-2025, suggesting improved security practices are making a meaningful difference.*

And then, to give put some more flesh on this, they explain:

*The cryptocurrency ecosystem faced another challenging year in 2025, with stolen funds continuing their upward trajectory. Our analysis reveals a shift in crypto theft patterns, characterized by four key developments: the persistence of the Democratic People's Republic of Korea (DPRK) as a primary threat actor, the growing severity of individual attacks on centralized services, a surge in personal wallet compromises, and an unexpected divergence in decentralized finance (DeFi) hack trends.*

*These patterns emerge clearly from the data and reveal significant changes in how crypto theft is occurring across different platform types and victim categories. As digital asset adoption expands and valuations reach new heights, understanding these evolving security threats has become increasingly critical.*

*The cryptocurrency industry witnessed over \$3.4 billion in theft from January through early December 2025, with the February compromise of Bybit alone accounting for \$1.5 billion of that total.*

We covered the record breaking Bybit hack back in early March of this year. Recall that they used a third party multisig wallet provider service called "Safe{Wallet}". The extremely clever North Korean hackers injected malicious code into the Safe{Wallet} domain which selectively targeted Bybit's smart contracts and multi-signature process.

But aside from all that, stop to consider that just during 2025, this year, the concerted efforts of North Korean hackers netted the DPRK \$3.4 billion US dollars in digital cryptocurrency. That's something.

Chainalysis continues:

*Beyond the headline figure, the data reveal important shifts in the composition of these thefts. Personal wallet compromises have grown substantially, increasing from just 7.3% of total stolen value in 2022 to 44% in 2024. In 2025, the share would have been 37% if it weren't for the outsized impact of the single Bybit attack.*

*Meanwhile, centralized services are experiencing increasingly large losses due to private key compromises. Despite their institutional resources and professional security teams, these platforms remain vulnerable because of this fundamental security challenge. While such compromises are infrequent, their scale still drives enormous shares of stolen volume when they do occur, accounting for 88% of losses in Q1 2025.*

*The persistence of high theft volume indicates that while some areas of crypto security may be improving, attackers continue to find success across multiple vectors.*

*Stolen fund activity has always been outlier-driven, with most hacks relatively small and some immense. But 2025 reveals a striking escalation: the ratio between the largest hack and median of all incidents has crossed the 1,000x threshold for the first time. Funds stolen in the largest attacks are now 1,000 times larger than those stolen in the typical incident, surpassing even the 2021 bull market peak. These calculations are based on the USD values of funds stolen at the time of their theft.*

*This growing discrepancy has concentrated losses dramatically. The top three hacks in 2025 account for 69% of all service losses, creating a landscape where individual incidents have an outsized impact on yearly totals. While the number of incidents may fluctuate and median losses grow with asset prices, the potential for catastrophic individual breaches is escalating faster still.*

*The Democratic People's Republic of Korea (DPRK) continues to pose the most significant nation-state threat to cryptocurrency security, achieving a record-breaking year for stolen funds despite an assessed dramatic reduction in attack frequency. In 2025, North Korean hackers stole at least \$2.02 billion in cryptocurrency (\$681 million more than 2024), representing a 51% increase year-over-year. This marks the most severe year on record for DPRK crypto theft in terms of value stolen, with DPRK attacks also accounting for a record 76% of all service compromises. **Overall, 2025's numbers bring the lower-bound cumulative estimate for cryptocurrency funds stolen by the DPRK to \$6.75 billion.***

*North Korean threat actors are increasingly achieving these outsized results often by embedding IT workers – one of DPRK's principal attack vectors – inside crypto services to gain privileged access and enable high-impact compromises. Part of this record year likely reflects an expanded reliance on IT worker infiltration at exchanges, custodians, and web3 firms, which can accelerate initial access and lateral movement ahead of large-scale theft.*

*More recently, however, DPRK-linked operators have flipped this IT worker model on its head. Instead of merely applying for roles and embedding themselves as employees, they are increasingly impersonating recruiters for prominent web3 and AI firms, orchestrating fake hiring processes that culminate in "technical screens" designed to harvest credentials, source code, and VPN or SSO access to the victim's **current** employer. At the executive level, a similar social-engineering playbook appears in the form of bogus outreach from purported strategic investors or acquirers, who use pitch meetings and pseudo-due diligence to probe for sensitive systems information and potential access paths into high-value infrastructure – an evolution that builds directly on the DPRK's IT worker fraud operations and their focus on strategically important AI and blockchain companies.*

*As we have seen in years past, the DPRK continues to undertake significantly higher-value attacks than other threat actors. As shown in the chart below, from 2022-2025, DPRK-attributed hacks occupy the highest value ranges, while non-DPRK hacks show more normal distributions across all theft sizes. This pattern reinforces that when North Korean hackers strike, they target large services and aim for maximum impact.*

*This year's record haul came from significantly fewer known incidents. This shift — fewer incidents yielding far greater returns — reflects the impact of the massive Bybit hack in February 2025.*

Another way to look at this is that they are leveraging trust at every level. They might observe and determine who provides janitorial services for an intended target, then obtain employment there and arrange to obtain access to their target's physical facilities. Or they masquerade as recruiters who are attempting to hire employees away from their target and use the departing employee's access and desire to switch up to a higher paying job to gain useful inside information. Unfortunately, while we're all hard at work generating income, they're spending their time coming up with new and clever ways to separate us from that income we've generated.

Chainalysis then spends some time talking about the money laundering employed by the DPRK. But their discussion of the escalating threat from the number, if not the size, of personal wallet compromises was interesting. They wrote:

*Through analysis of on-chain patterns, in addition to the reporting from victims and industry partners, we can gain an understanding of the magnitude of personal wallet compromises, although the true number of compromises is likely far greater.*

*Based on our lower bound estimates, personal wallet compromises now account for 20% of all value stolen in 2025, down from 44% of the total in 2024, representing an evolution in both scale and pattern. The total number of theft incidents surged to 158,000 in 2025, nearly triple the 54,000 recorded in 2022. Unique victims increased from 40,000 in 2022 to at least 80,000 in 2025. These dramatic increases are likely due to greater crypto adoption. For example, Solana, one of the blockchains with the greatest number of active personal wallets, had by far the largest number of incidents (~26,500 victims).*

*Yet despite more incidents and victims, the total US Dollar value stolen from individual victims actually declined from 2024's peak of \$1.5 billion to \$713 million in 2025. This suggests that attackers are targeting more users, but are managing to steal smaller amounts per victim.*

*Network-specific victimization data provides additional insight into which domains present the greatest risk to crypto users. When measuring crime rates per 100K wallets in 2025, Ethereum and Tron show the highest rates of theft. Ethereum's large size indicates both high rates of theft and high victim count, while Tron's position shows elevated rate of theft despite a smaller active wallet base. In contrast, Base and Solana show lower victimization rates despite significant user bases.*

*These measurable differences highlight that personal wallet security risks are not uniform across the crypto ecosystem. The variation in victimization rates across chains with similar technical architectures suggests that factors beyond technology — such as user demographics,*

*popular applications, and criminal infrastructure – play important roles in determining theft rates.*

What all this says, ultimately, for the end user is that anyone who is technically capable of transferring any cryptocurrency they do not need to have online into an offline wallet has nothing to lose and everything to gain. If I owned any appreciable amount of cryptocurrency I would not be inclined to leave it sitting in an online account. The beauty of this technology is that another wallet can be created with a private key that has never been online, and the currency can be safely transferred there. It's true that you are then utterly responsible for its safe keeping – which is a lesson that Leo and I both painfully learned the hard way back before our crypto had become valuable.

### **Cryptomining discovery by AWS abusing EC2 and ECS (and Docker)**

Last week, Amazon's AWS Security Blog shared the news of their discovery of an advanced cryptomining operation targeting AWS users whose credentials had leaked. The brief start of their blog posting reads like a sales and market piece, but I'll share it as a means of establishing the context. They wrote:

*Amazon GuardDuty and our automated security monitoring systems identified an ongoing cryptocurrency (crypto) mining campaign beginning on November 2, 2025. The operation uses compromised AWS Identity and Access Management (IAM) credentials to target Amazon Elastic Container Service (ECS) and Amazon Elastic Compute Cloud (EC2). GuardDuty Extended Threat Detection was able to correlate signals across these data sources to raise a critical severity attack sequence finding. Using the massive, advanced threat intelligence capability and existing detection mechanisms of Amazon Web Services (AWS), GuardDuty proactively identified this ongoing campaign and quickly alerted customers to the threat. AWS is sharing relevant findings and mitigation guidance to help customers take appropriate action on this ongoing campaign.*

*It's important to note that these actions don't take advantage of a vulnerability within an AWS service, but rather require valid credentials that an unauthorized user uses in an unintended way. Although these actions occur in the customer domain of the shared responsibility model, AWS recommends steps that customers can use to detect, prevent, or reduce the impact of such activity.*

Okay. So, essentially, our GuardDog sniffed out some suspicious activity (oh, by the way, using its massive threat intelligence) and we found that bad guys were abusing our customers' accounts after having somehow obtained their IAM account credentials. Then we begin to get some additional interesting details. They write:

*The recently detected crypto mining campaign employed a novel persistence technique designed to disrupt incident response and extend mining operations. The ongoing campaign was originally identified when GuardDuty security engineers discovered similar attack techniques being used across multiple AWS customer accounts, indicating a coordinated campaign targeting customers using compromised IAM credentials. Operating from an external hosting provider, the threat actor quickly enumerated Amazon EC2*

*service quotas and IAM permissions before deploying crypto mining resources across Amazon EC2 and Amazon ECS. Within 10 minutes of the threat actor gaining initial access, crypto miners were operational.*

*A key technique observed in this attack was the use of "ModifyInstanceAttribute" with "Disable API Termination" set to TRUE, forcing victims to re-enable API termination before deleting the impacted resources. Disabling instance termination protection adds an additional consideration for incident responders and can disrupt automated remediation controls. The threat actor's scripted use of multiple compute services, in combination with emerging persistence techniques, represents an advancement in crypto mining persistence methodologies that security teams should be aware of.*

This use of "Disable API Termination" (also known as termination protection) is a setting on an Amazon EC2 instance that prevents that instance from being terminated using AWS-provided APIs, the AWS command-line interface, or the AWS Management Console. Clearly, the intent is to give EC2 users a means of preventing the accidental termination of some service or process that absolutely needs to always be present and running. So it's not surprising that bad guys who know their way around the operation of AWS EC2 compute services would enable blocking their cryptominer's termination.

We then learn that a malicious Docker Hub image was created a few days prior to the first observed instance of this intrusion, on October 29th, with over 100,000 pulls. And that Docker Hub image was used to deploy crypto miners to containerized environments. Inside that image they found an SRBMiner-MULTI binary for crypto mining. This specific image, having been identified as malicious, has since been taken down from Docker Hub, but we know that threat actors will probably deploy similar images under different names.

The AWS security guys also discovered that the attackers employed the AWS SDK for Python (Boto3) user agent to deploy Python-based automation scripts throughout the entire attack chain. Crypto mining domains: asia[.]rplant[.]xyz, eu[.]rplant[.]xyz, and na[.]rplant[.]xyz were used.

Amazon's mention of the SRBMiner reminded me of something I had seen earlier. So I tracked down a recent piece in The Hacker News titled "Cybercriminals Exploiting Docker API Servers for SRBMiner Crypto Mining Attacks". The Hacker News wrote:

*Bad actors have been observed targeting Docker remote API servers to deploy the SRBMiner crypto miner on compromised instances, according to new findings from Trend Micro. The Trend Micro researcher said: "In this attack, the threat actor used the gRPC protocol over h2c to evade security solutions and execute their crypto mining operations on the Docker host. The attacker first checks the availability and version of the Docker API, then proceeds with requests for gRPC/h2c upgrades and gRPC methods to manipulate Docker functionalities."*

*The adversary checks for gRPC methods that are designed to carry out various tasks pertaining to managing and operating Docker environments, including those related to health checks, file synchronization, authentication, secrets management, and SSH forwarding.*

There's a bunch of Docker-specific jargon flying back and forth here. So what's happening is that the Internet now contains a population of publicly accessible Docker remote API servers which, when not properly secured, can be remotely exploited to accept, host and run attacker-provided cryptominers.

The Docker Docs talk about this. There's a page titled "[Configure remote access for Docker daemon](#)" which says:

*By default, the Docker daemon listens for connections on a Unix socket to accept requests from **local** clients. You can configure Docker to accept requests from remote clients by configuring it to listen on an IP address and port as well as the Unix socket.*

But then, in a big impossible-to-miss warning box, the page says:

**WARNING!!** — *Configuring Docker to accept connections from remote clients can leave you vulnerable to unauthorized access to the host and other attacks. It's critically important that you understand the security implications of opening Docker to the network. If steps are not taken to secure the connection, it's possible for remote non-root users to gain root access on the host. Remote access without TLS is not recommended, and will require explicit opt-in in a future release. For more information on how to use TLS certificates to secure this connection, see [Protect the Docker daemon socket](#).*

So we learn that Docker, themselves, did everything right. The default is secure, local machine-only, access by clients running on the local machine, connecting to Docker through the local Unix socket interface. So, it appears that there are those who wanted to have their Docker instances available across the network. Did they intend it only for the LAN and not the WAN? Was this a misconfiguration of an important option? Or did they deliberately make their Docker instances available across the entire global Internet?

I should spend some time distilling a short list of fundamental laws of security. Issac Asimov created his three laws of robotics. This podcast could have a similar short set of laws. If we did have such, up there right near the top would be: "*Never rely upon the strength of remote authentication.*" That would have to be one of the golden rules. We keep seeing that mistake being made over and over and over.

But for whatever reason, Docker's API is being published on the Internet and bad guys are now scouting around looking for them. So this is a variant on the AWS EC2 case we first talked about. In this second instance, bad guys have figured out a way to bypass several layers of intended security. Trend Micro and The Hacker News both concluded their coverage with the advice to better secure all publicly exposed instances of Docker API servers. Yeah.

I should also clarify that the SRBMiner that was implicated in both of these cases is not in any way malicious itself. It's a beautiful piece of work. It's a CPU + GPU miner which mines using a system's processor plus AMD, NVIDIA, or Intel GPUs depending upon the build. It's able to mine using up to four different algorithms, which is to say types of cryptocurrency, simultaneously which is why it's called SRBMiner-Multi. It's available to run on 64-bit instances of either Windows or Linux, and it can be found at [srbminer.com](http://srbminer.com).

Poking around over there we see a list of interesting features:

- Mine up to 4 algorithms simultaneously
- Guided setup mode
- Run in background without a window
- Hashrate watchdog that restarts miner on GPU error
- Monitoring of GPU temperature, and auto turn off if temperature is too high
- System shutdown on too high GPU temperature
- Miner auto restart on too many rejected shares
- API for miner statistics
- Web based GUI interface for miner statistics
- Multiple pools with failover support
- Difficulty monitor, reconnects to pool if difficulty is too high
- Job timeout monitor, reconnects to pool if no job received for a long time
- Bunch of other useful features

SRBMiner-Multi is distributed there, at the project's official site and also via a GitHub repo.

The reason I wanted to share these two recent examples of surreptitious mining is that they dovetail so nicely with the Chainalysis report about North Korea. **All** of these instances have a single common thread. And that thread — is MONEY. It's about money. It's all about money and it's only about money.

That's also, of course, the entire motivating factor behind all of the breaches and ransomware and extortion. The bad guys want to obtain an advantage. And they want to leverage that advantage to get themselves as much of someone else's money as they can. They could not care less about some random company's client list or random people's social security numbers or anything else that might be stored in an exfiltrated database. But if they can figure out a way to turn that data — which they themselves have no interest in whatsoever — into some cold hard cash, then, unfortunately for the original owners of that data, they will be highly motivated to do just that. So it's all about money. They want ours. And, sadly, today's network and other security practices are proving not to be strong enough to keep them from finding ways to get it.

### **Why's our TV being so sluggish?**

I ran across an interesting description of a new, quite large and capable Android-based DDoS Botnet that preferentially inhabits SmartTVs. This Botnet appears to be capable of generating around 30 terabits per second of DDoS flooding traffic and it also has many other features that would concern anyone who knew that it had taken up residence in their family's SmartTV.

The security company that received a sample of this Bot and reverse-engineered its operation posted their complete analysis under the title: *"Kimwolf Exposed: The Massive Android Botnet with 1.8 Million Infected Devices"*. They wrote:

*On October 24, 2025, a trusted partner in the security community provided us with a brand-new botnet sample. The most distinctive feature of this sample was its C2 domain, 14emeliaterracewestroxburyma02132[.]su, which at the time ranked 2nd in the Cloudflare*

*Domain Rankings. A week later, it even surpassed Google to claim the number one spot in Cloudflare's global domain popularity rankings. There is no doubt that this is a hyper-scale botnet. Based on the information output during runtime and its use of the wolfSSL library, we have named it Kimwolf.*

Just to clarify what they intend by citing Cloudflare's Domain Rankings, Cloudflare tracks, ranks and reports the domains being used across the Internet. There are so many instances of this newly discovered Botnet that it was briefly taking the #1 slot in Cloudflare's rankings, pushing even Google down into second place. That's a lot of activity.

Get a load of what they have discovered about this massive newcomer:

*Kimwolf is a botnet compiled using the NDK (Android's Native Developer Kit). In addition to typical DDoS attack capabilities, it integrates proxy forwarding, reverse shell, and file management functions. From an overall architectural perspective, its functional design is not complex, but there are some highlights worth noting: for example, the sample uses a simple yet effective Stack XOR operation to encrypt sensitive data; meanwhile, it utilizes the DNS over TLS (DoT) protocol to encapsulate DNS requests to evade traditional security detection.*

*Furthermore, its Command & Control identity authentication employs a digital signature protection mechanism based on elliptic curves, where the Bot side will only accept communication instructions after the signature verification passes. Recently, it has even introduced EtherHiding technology to counter takedowns using blockchain domains. These features are relatively rare in similar malware. Based on our analysis results, it primarily targets Android platform TV boxes. The "Welcome to Android Support Center" message displayed on the Command & Control backend also corroborates this.*

*The Kimwolf samples use a naming rule to identify version numbers. The sample previously provided by our community partner was version v4. After completing the reverse engineering analysis, we imported the sample's intelligence into the XLab's Cyber Threat Insight and Analysis System, successively capturing multiple related samples including v4 and v5, thus achieving automated continuous tracking of this family.*

*On November 30, we captured another new sample of this botnet family and successfully took over one of the C2 domains, thereby obtaining the opportunity to directly observe the true operating scale of this botnet for the first time. Based on statistics from source IP data that established connections with our registered C2 address and whose communication behavior matched Kimwolf C2 protocol characteristics, we observed a cumulative total of approximately 2.7 million distinct source IP addresses over the three days from December 3 to December 5.*

*Among them, we observed approximately 1.36 million active IPs on December 3, about 1.83 million on December 4, and about 1.5 million on December 5 (there is IP overlap between different dates). Analysis indicates that Kimwolf's primary infection targets are TV boxes deployed in residential network environments. Since residential networks usually adopt dynamic IP allocation mechanisms, the public IPs of devices change over time, so the true scale of infected devices cannot be accurately measured solely by the quantity of IPs. In other words, the cumulative observation of 2.7 million IP addresses does not equate to 2.7 million infected devices.*

*Despite this, we still have sufficient reason to believe that the actual number of devices infected by Kimwolf exceeds 1.8 million. This judgment is based on observations in the following areas:*

- Kimwolf uses multiple C2 infrastructures. We took over only a portion of the C2s, so we could only observe the activity of some Bots, unable to cover the full picture of the botnet.*
- On December 4, the number of Bot IPs we observed reached approximately 1.83 million, a historical peak. On that day, parts of the C2s normally used by Kimwolf were taken down by relevant organizations, causing a large number of Bots to fail to connect to the original C2s and turn to try connecting to the C2 we preemptively registered. This anomalous event caused more Bots to be centrally exposed in a short period, so the data for that day may be closer to the lower limit of the true infection scale.*
- Infected devices are distributed across multiple global time zones. Affected by time zone differences and usage habits (e.g., turning off devices at night, not using TV boxes during holidays, etc.), these devices are not online simultaneously, further increasing the difficulty of comprehensive observation through a single time window.*
- Kimwolf exists in multiple different versions, and the C2s used by different versions are not completely identical, which is also one of the important reasons why we cannot obtain a complete perspective.*

*Combining the above factors, we conservatively estimate that the actual number of devices infected by Kimwolf has exceeded 1.8 million. A botnet of such scale possesses the capability to launch massive cyberattacks, and its potential destructive power cannot be ignored.*

*While working hard to track new versions, we were also full of curiosity about the old versions. Through source tracing analysis, although we failed to capture old versions like v1 or v2, we surprisingly found that Kimwolf is actually associated with the Aisuru botnet. Kimwolf relies on an APK file to load and start it during runtime. A DEX file uploaded to VT from India on October 7 showed obvious homologous characteristics with Kimwolf's APK. Subsequently, on October 18, the parent APK of that DEX was uploaded to VT from Algeria; the resource files of this APK contained Aisuru samples for 3 CPU architectures: x86, x64, and arm. We speculate that in the early stages of this campaign, the attackers directly reused Aisuru's code; subsequently, likely because Aisuru samples had high detection rates in security products—Android platforms have more mature security protection systems compared to IoT ecosystems—the group decided to redesign and develop the Kimwolf botnet to enhance stealth and evade detection.*

*From the monitoring data of the XLab command tracking system, statistics show that the main functions of the Kimwolf botnet are usually concentrated on traffic proxying, with a small amount of DDoS attacks. However, between November 19 and 22, it suddenly went "crazy": in just 3 days, it issued 1.7 billion DDoS attack commands, with the attack range covering massive amounts of IP addresses globally. This high-profile spree follows on the heels of the C2 domain's unprecedented rise to #1 in global popularity. Theoretically, such a large number of attack commands and targets may not be able to produce substantial attack effects on the targets; this behavior may have been purely to demonstrate its own presence.*

*Currently, the security community's understanding of Kimwolf presents a polarized situation. Information in the public intelligence field is scarce, its propagation path is not yet clear, and the detection rate of related samples and their C2 domains on VirusTotal is extremely low. At the same time, due to the adoption of covert technologies like (DoT), the association between*

*its C2 and samples has not been effectively discovered. However, at the non-public threat confrontation level, the situation is entirely different. We observed that Kimwolf's C2 domains have been successfully taken down by unknown parties at least three times, forcing it to upgrade its tactics and turn to using ENS (Ethereum Name Service) to harden its infrastructure, demonstrating its powerful evolutionary capability. Given that Kimwolf has formed a massive attack scale, and its recent activity frequency and attack behaviors show a significant upward trend, we believe it is necessary to break the intelligence silence. We hereby release this technical analysis report to make relevant research results public, aiming to promote threat intelligence sharing, gather community strength to jointly respond to such threats, and effectively maintain cyberspace security.*

Okay. Now everyone has a good sense for what's going on with this apparent descendent of the previously massive and famous Aisuru botnet.

So, one question is, where are these infected consumer TV boxes? Since these researchers were briefly in the position to be receiving incoming Bot traffic to their command and control IP, they were able to obtain the bot's demographics: Infected devices are distributed in 222 countries and regions globally. The top 15 countries are analyzed as: Brazil 14.63%, India 12.71%, USA 9.58%, Argentina 7.19%, South Africa 3.85%, Philippines 3.58%, Mexico 3.07%, China 3.04%, Thailand 2.46%, Saudi Arabia 2.37%, Indonesia 1.87%, Morocco 1.85%, Turkey 1.60%, Iraq 1.53%, Pakistan 1.39%.

I'll share one more piece from their extensive research. They wrote:

*Readers familiar with DDoS might be curious: "For such a huge botnet, what level has its attack capability actually reached?" Although we cannot directly measure it, through observations of two large-scale DDoS events and a horizontal comparison with Aisuru, we believe Kimwolf's attack capability is close to 30Tbps.*

- A well-known cloud service provider observed a 2.3Bbps attack at 22:09Z on November 23, with 450,000 participating IPs. We confirmed Kimwolf's participation.*
- A well-known cloud service provider observed an attack nearing 30Tbps and 2.9Gpps at 09:35Z on December 9. After data comparison, both parties confirmed Kimwolf's participation.*
- Cloudflare pointed out in its 3rd quarter 2025 DDoS threat report that Aisuru is one of the strongest known botnets currently, with a control scale of millions of IoT/network devices, capable of sustaining Tbps-level attacks and even peak attacks approaching 30 Tbps and more than 10 Bpps.*

*In fact, we believe that behind many attacks observed by Cloudflare attributed to Aisuru, it may not just be the Aisuru botnet acting alone; Kimwolf may also be participating, or even led by Kimwolf. These two major botnets propagated through the same infection scripts between September and November, coexisting in the same batch of devices. They actually belong to the same hacker group.*

If 9.58% of Kimwolf infections have been seen in the US, and if there are conservatively more than 1.8 million operating instances of Kimwolf, that's more than 172,000 Android-based SmartTVs currently infected with Kimwolf in the United States.

They conclude their very thorough analysis, writing:

*This is the majority of the intelligence we currently possess on the Kimwolf botnet. Giant botnets originated with Mirai in 2016, with infection targets mainly concentrated on IoT devices like home broadband routers and cameras. However, in recent years, information on multiple million-level giant botnets like Badbox, Bigpanzi, Vo1d, and Kimwolf has been disclosed, indicating that some attackers have started to turn their attention to various smart TVs and TV boxes. These devices generally suffer from problems like firmware vulnerabilities, pre-installed malicious components, weak passwords, and lack of security update mechanisms, making them extremely easy for attackers to control long-term and use for large-scale cyberattacks. One of our motives for disclosing the Kimwolf botnet this time is to call on the security community to give due attention to smart TV-related devices.*

*After attackers gain root privileges on smart TVs, the resulting attacks are not limited to traditional cyberspace. Attackers can use controlled terminals to insert tampered, biased, or extreme videos. In the legal systems of many countries, inserting content without written permission violates the contract between the viewer and the TV program provider and is illegal. This is our second motive for disclosing the Kimwolf botnet this time, calling on law enforcement agencies to consider scrutinizing such suspected illegal activities related to smart TVs.*

*Against the backdrop of overlapping threats, whether ordinary TV box users, sales channels, operators, or regulatory departments and manufacturers, all must attach great importance to the security of TV boxes. Among them, TV box users should especially: ensure devices come from reliable sources, use firmware that can be updated in time, avoid setting weak passwords, and refuse to install APKs of unknown origin to reduce the risk of being infected and controlled by botnets.*

*We sincerely welcome CERTs from all countries to contact us, share intelligence and vision, join hands to combat cybercrime, and jointly maintain global cybersecurity. If you are interested in our research, or have inside information, feel free to contact us via X platform.*

I have placed a link to their entire analysis, most of which I skipped over because it's more than we need here. But for anyone who wants to get a very clear look into the guts of a massively successful state-of-the-art global Botnet, these guys have published that:

<https://blog.xlab.qianxin.com/kimwolf-botnet-en/>

## DNS Benchmark Update

I'm very pleased with the commercial launch of version 2 of the DNS Benchmark after a year of work. It's still in the process of settling down and is now in its 4th release. It acquired a couple of new features, and fixed two bugs that had escaped notice until now. Windows 11's new "Smart App Control" blocked another person's use of the product, but now we knew to ask them to try again. When they did, they had no trouble. So far, no one has been permanently blocked. So GRC has a new solid commercial offering, and a smooth updating process that gives us the ability to move the product forward to add features and fix bugs that may arise.

# Listener Feedback

## Jamie

*Hello Steve, huge fan and very long time listener. Just wanted to give you some quick information that might be helpful to your listeners.*

*A very quick and painless way to run the DNS benchmark on any linux system is to install Steam, add a non-steam executable to your library and use Proton as the compatibility layer, it takes about 15 seconds and the benchmark runs perfectly!*

*And... A couple of episodes you mentioned wanting some more insights into traffic entering and exiting your network. Take a look at the Netdata plugin for pfsense. If I have any bead on your interests, you might want to set aside an afternoon to dive into it. It gives you an incredibly deep and insightful look into your traffic.*

*Thanks for everything you do, much love to you and Leo! Thanks, Jamie in Las Vegas...*

This turns out to be a terrific suggestion. I wonder whether it also works for Mac. As I understand it, the step-by-step is:

1. Install Steam on Linux
2. Launch Steam → Add a Game → Add a Non-Steam Game
3. Select the Windows DNS benchmark .exe
4. Right-click the entry → Properties
5. Enable "Force the use of a specific compatibility tool"
6. Select Proton (latest stable)
7. Click Play
8. That's it.

I've been looking for a simple way to solve the "Running GRC's Windows Apps on Linux and Mac" problem for people who may not have needed to do so before, so this is a great tip. THANK YOU Jamie! As we know, the newly commercial DNS Benchmark is just the first of several planned Windows apps I hope to produce, so having a solid way to make that happen will be terrific.

## Rick Andrews

*Steve, in this episode you noted that the hundreds of millions of certificates issued by Let's Encrypt represented a huge risk, saying that: "a billion websites are all now dependent upon a single service for their certificates ". But many other public CAs, including DigiCert, offer ACME-based service to automatically obtain a certificate that chains up to one of their roots. In other words, you can use ACME with someone other than Let's Encrypt, and if more people did that, it would reduce and spread out the risk. I just wanted to clarify that. -Rick Andrews*

Rick is absolutely correct. But to the best of my knowledge there are only two providers of Domain Validation (DV) web certificates who offer them at no charge. That's Let's Encrypt and ZeroSSL. But ZeroSSL wants to sell you stuff. And they show that their free certs are limited to 3 per customer. They also require you to create an account, verify your email and all that rigmarole we're all too familiar with. So, in my opinion, there's only Let's Encrypt who has the fundamentally correct ethics about truly free TLS certificates.

Looking at the point Rick makes another way, it's utterly obvious that with the shortening

life-cycle of TLS web certificates eventually marching down to 47 days, ANY certificate authority that wishes to remain in business must already have, or be rapidly working to, bring ACME certificate issuance automation online.

Thinking about this caused me to wonder who exactly IS paying the bills for Let's Encrypt, a service that's wonderful to have — thanks very much! — but is also quite easily taken for granted. Set it up and forget it. Problem solved. But a number of times we've looked at the scaling that Let's Encrypt needs to do, especially as certificate lifetimes continue to shorten. So, again, is this a truly free lunch?

After a bit I digging, here's what I found:

As I noted last week, Let's Encrypt is operated by and a service of the nonprofit Internet Security Research Group – ISRG. And the ISRG is funded entirely through charitable contributions, sponsorships, grants, and donations from individuals and organizations that support its mission to encrypt the web. 100% of its funding comes from these contributed sources rather than user fees. So who's known to be providing this support? Major organizational supporters & financial sponsors include Google, the Mozilla Foundation, Cisco, OVHcloud, Facebook/Meta, AWS, Shopify, Nginx, the Internet Society (ISOC), SiteGround, Automattic, Hostpoint, Discourse, infomaniak and PlanetHoster. Additionally, the EFF and the Ford Foundation are backers and the Open Technology Fund (OTF) has been reported to provide grant support to ISRG.

I've never stopped to think about the question of who pays for all this from Let's Encrypt. I'm not yet using Let's Encrypt certs, I'm still happily with DigiCert. But the decisions the CA/Browser forum have made regarding web certificate lifetimes means that I'll be moving to Let's Encrypt and also that I plan to be voluntarily supporting them, much the way I do Wikipedia, because having access to Let's Encrypt is a privilege that should never be taken for granted.

This brought me to wonder about the stance of a major – the major – certificate authority with whom I have proudly hung my hat since I left Verisign. As a DigiCert customer I've received their email announcing their support for ACME certificate issuance automation. I like DigiCert, so I wanted to be certain that I would not be able to remain with them. I went over to DigiCert and used their site-search to search for "*Free SSL/TLS Certificates*". The first link that came up was titled: "*The Fraud Problem with Free SSL Certificates*" — I thought: "*Oh! ... THIS ought to be interesting. What does the company that's never been in the business of issuing free certificates have to say about those who do?*" Here's DigiCert's take on why they do not offer free web certificates:

*SSL Certificates are the defacto standard for online trust today. SSLs are so important to online security that Google gives a ranking boost to sites that secure their content with HTTPS.*

If, like me, at this point you're thinking "What??!" You cannot have a site today that's not HTTPS. And I was already curious due to their use of SSL instead of TLS. So I went looking for a date on this posting and I found it: April 6th, 2015. So this is ten and a half years ago. I still wanted to know what they thought, especially since their policy hasn't appeared to change. And I wanted to see whether there might be anything to learn. They wrote:

*Savvy Internet users have come to recognize and expect that any website asking for sensitive or personal information to display the universal symbol—the padlock—before typing in any sensitive information. In a Tech-Ed survey, users reported that without knowing the identity of*

*the organization conducting business, over 35% would reconsider entering a credit card number from a site using a plain SSL Certificate.*

*Are SSLs less trustworthy than we think? To answer this question, we have to consider the fact that not all SSL Certificates are created equal.*

**Domain Validated (DV):** *No identity verification is done. The Certificate Authority (CA) sends an automated challenge email and the site owner clicks on a link to approve the certificate. Information is encrypted, but no assurance is made that the organization should be trusted. Because of the lack of trust and the frequent use for fraudulent purposes, **DigiCert does not issue cheap domain validated certificates.***

Right. And they didn't say that DigiCert did not issue "free" domain validated certificates. They said they didn't issue "cheap" domain validated certificates. What do they have to say about the other sorts of certs?

**Organization Validated (OV):** *Basic identity verification is completed. In the case of an OV certificate, the CA conducts a much more substantial validation process. This includes checking the applicant's business credentials (through government and business databases) and verifying that the website is a legitimate organization. DigiCert validation experts are online 24/7 and can complete basic verification in less than 10 minutes on most certificates.*

**Extended Validation (EV):** *Extended identity verification is completed. This is the highest level of validation and strict standards for identity verification. The validation process includes physical location checks, phone calls to ensure the applicant is authorized to order the certificate on behalf of the company or business represented, and more. DigiCert EV is issued in less than 24 hours for most EV Certificate requests.*

*Although all SSLs ensure that information online is encrypted, only OV and EV SSL Certificates actually certify the website is being operated by a legitimate organization, keeping users safe from fraud and phishing scams online.*

### **The Problem with Free SSL Certificates**

*Let's be honest, no one can give something away for free and remain in business for very long. Some organizations today provide free SSL Certificates, relying solely on automated systems that skip authentication to keep costs extremely low. These organizations often provide add-on services for a fee, or are funded by third-party organizations with deep pockets.*

*Authentication is critical to online trust. Authentication provides the assurance that you're at the real PayPal.com, and not a fake PayPal phishing site. CAs that include identity authentication in their certificates follow strict rules for verifying identity of organizations, individuals, and the authority to request SSL Certificates on behalf of organization. Free SSL Certificates don't rely on performing authentication checks or identify verification, making them a prime candidate for fraudulent websites today.*

*Jerome Segura of Malwarebytes reported on an email campaign that leveraged a site benefiting from a free CloudFlare certificate in order to deliver malware to users online.*

Oh, right! I'd forgotten about Cloudflare! While Cloudflare is not an ACME user, anyone using Cloudflare's hosting – including their free tier – gets HTTPS connections at no cost. So a website

with an SSL certificate is delivering malware online. Nothing about DV vs OV or EV would prevent that. The truth is, users don't know what sort of certificate a site is using, they don't understand the distinctions and they don't care. DigiCert's posting continues:

*The malicious email message claimed to be a notice from cloud-based, remote connectivity service provider LogMeIn, about an alleged problem with extending the service subscription due to insufficient funds. The HTTPS link included in the email claimed to point to an invoice showing the details of the transaction. Since the website had an SSL Certificate installed, users were more likely to trust it and download the malware file.*

I doubt that was true even 10 years ago, and it's certainly not true now, since all websites must use encryption for today's browsers to have anything to do with them. They continue:

*Fortunately, CloudFlare has since revoked the certificate for the website and the location is now flagged as malicious in all major web browsers.*

*However, this is only the tip of the iceberg and cyber criminals are taking notice. With free SSL Certificates or cheap SSL becoming more readily available, it's likely that cyber criminals will continue to exploit the lack of identity verification to take advantage of users online.*

They then switch into their "EV Certificates for All" pitch and while much about that has changed in the last 10 years, I want to share that, too, because I have something specific to share about that. DigiCert wrote:

*In working with the CA/Browser Forum industry group to create EV SSL Certificates, DigiCert set out to ensure that any organization could qualify for an EV certificate. We continue to work with the group to make amendments to the EV verification process to ensure that more organizations can take advantage of the higher level of trust the EV provides, while ensuring that the process remains cyber-crime proof. Tech-Ed's EV survey showed that 67% of web users said they would not buy from an unfamiliar website that did not have an EV SSL Certificate to confirm the identity of the organization.*

Really? Does anyone believe that? Even then years ago when the presence of EV certs was apparent, I find it difficult to believe that 2/3rds of random buyers would have lifted their hand from their mouse and backed away from their keyboard upon seeing a lack of Extended Validation. Again, no one ever really knew what that was about so the browsers all removed it. DigiCert concludes:

*Microsoft even adopted EV as their code signing standard for application security and require all UEFI code submissions must be signed by an Extended Validation (EV) Code Signing Certificate.*

And we know that has also changed. There is no longer any extra benefit from code signing with EV. They said:

*Enterprise Benefits of Extended Validation: EV SSL Certificates ensure that users can communicate securely with a website. Websites using an EV SSL Certificate gain immediate trust in the eyes of users because it reassures the user that the data is secure and the*

*organization receiving the data is a reputable entity. Since technical requirements prevent EV SSL Certificates from being forged, large enterprises especially benefit from using EV certificates as an easy anti-phishing indicator or that data being secured cannot be intercepted by a malicious third party. Keeping users safe online and staying ahead of cyber criminals and scammers requires going above and beyond in online security. Identity verification is the clear answer to the problem of online trust.*

Okay. If I were attempting to find some valid contemporary benefit to justify the investment in Extended Validation certificates today, it would be machine-to-machine API-style security. With organizations all being linked up with outsourced business processes, I can see that requiring an inter-machine link to be EV validated could make sense.

With automated systems, you don't need to worry about domain name typos or phishing with look-alike domain names, automated domain names will not be mistyped. So strengthening that form of automated security could make sense.

But I cannot see ANY – not **any** – user-facing instance today where the type of TLS certificate makes any difference. It's disappeared from the browser's UI and no one is going to go digging down to examine the details of a website's certificate.

The personal experience I wanted to share regarding Extended Validation was one I had just last week with DigiCert. Our listeners who have been with the podcast for a while will recall the adventure of me establishing a contemporary email system for GRC. As part of that I wanted to experiment with BIMI – the system that is supposed to increase email trust by having GRC's logo able to appear in the recipient's inbox for those providers who support it.

Wouldn't you know it? BIMI certificates require full Extended Validation assurance. And like all other certificates, my original BIMI certificate was nearing its expiration date. Unfortunately, GRC's Extended Validation status, which also requires periodic renewal, had expired. And getting that done was such a hassle that I'm now glad there's no longer any benefit accruing from either EV code signing or EV web certificates.

To get re-EV certified, Sue needed to make an appointment to be present at our official corporate phone number published by Dun & Bradstreet or some other source of corporate intelligence. And appointments were booking four days out. Once that was done, I needed to engage in a video conference similar to the previous one, where I first sent DigiCert a high-res photo of my driver's license, then in front of a camera I followed instructions to look into the camera, then hold up the same driver's license, then move the license forward and pass my other hand back and forth in front of my face and behind the license. The very nice and very patient young woman on the other end of the zoom call, who had her camera on this time and I could see, explained that the hand waving was to prevent any sort of green screen from being used.

If Extended Validation still provided some actual value for web signing, or code signing, or even email signing, I'd be inclined to continue with it. But it's such an annoyance that I'm glad there's no longer any obvious value to be had. Since I've already purchased a 3-year BIMI certificate, I'll jump through whatever hoops I need to until those 3-years have expired. By then I'll hope that GRC will have established itself well enough as a known email provider that I can go BIMI-free. I get it that the industry wants the use of BIMI to actually mean something, but having it on every single piece of GRC's email from the start did not earn GRC any get out of jail free card. I still needed to battle the spam gods.

So returning to Rick's original point, where he wrote: *"But many other public CAs, including DigiCert, offer ACME-based service to automatically obtain a certificate that chains up to one of their roots."* Rick is 100% correct. But it appears that Let's Encrypt is in the unique position of having a business model that allows them to offer hassle-free, automatically issued & re-issued TLS web certificates. And now that they exist and are providing web certificates for nearly 2/3rds of the Internet they, along with Cloudflare, allow any and all websites to have TLS for free. So there's no way any other certificate authority whose business model requires revenue is ever going to compete with them.

Thanks, Rick, for a terrific discussion point!

**Jason Townsend** - reminds us of an old saying that's, sadly, less and less true today:

*Back in the 90's in a UK computer magazine there was a picture of a dog using a computer. The caption was 'The best thing about the Internet is that no one knows that you are really a dog'. Sadly it's getting more and more difficult to be a dog or a kangaroo on the internet and the days of anonymity are fading fast.*

Jason is referring to the famous New Yorker magazine cartoon published on July 5th, 1993:



**Jeff Root**

*Steve; Australia has done us a service, in that we now have great discussions about an important topic. Your piece was great, and got even better when Leo weighed in. But I think your insistence that age verification be "privacy protecting" is wrong-headed. Assume Apple and Google solve this problem perfectly. Now you go to a web site, maybe an online liquor store, and they use the age verification system which reveals nothing other than you are above a certain age. Now what?*

*Now you are let into a web site which is chock full of Google Analytics, Cloudflare Analytics, probably fronted by Cloudflare, and containing Javascript code from 30-50 other random sites. Security Now! has long reported on how easily ISPs and data brokers can de-anonymize users. So where's the privacy? And how has that effort to produce a 100% private age verification system made it harder for sites and data brokers to identify you?*

*I would suggest that Leo was right: the answer is not an Apple App, the answer is regulation and enforcement. Data brokers should be tightly regulated. Sites should be required to collect only such information as is necessary to render their pages, or transact business. What we need is a fully private and anonymous Internet, not yet another app which gives the illusion of*

*privacy. Just my opinion. Keep up the good work! Jeff Root / San Diego, CA*

I think Jeff makes a valid and ironic point about the idea of preserving privacy while gaining entry into a website where forces that are often beyond any visitor's control are all about tracking and profiling and doing everything possible to dissolve whatever privacy its users may imagine they have.

As for websites not collecting any information beyond that which is required to render their pages, the only way I can see that happening would be if the EU were to make that a requirement, much as they did with the cookie disclosure and permission pop-ups that they have made the entire world endure. They really did manage to change website behavior – for the worse. But I doubt we'll see the EU enforcing website privacy since European advertising and tracking companies are profiting just as much as companies everywhere else from the lack of effective privacy. We watched them kill Google's hopes for the Privacy Sandbox initiative that would have allowed interest profiling without cookies or tracking.

# GhostPoster

## How a PNG Icon Infected 50,000 Firefox Users

KOI calls themselves an endpoint security company. Last Tuesday, they published a nice descriptive piece about a recent discovery of theirs that immediately caught my eye and imagination, as I imagine it might catch our listeners'. So it what I wanted to share as this week's main topic. Under the headline: "Inside GhostPoster: How a PNG Icon Infected 50,000 Firefox Users", they explained:

*Every extension has a logo. A tiny image sitting in your toolbar, a visual shorthand for trust. You glance at it, you recognize it, you move on. You probably never think about what's actually inside that file. The authors of GhostPoster are counting on that.*



*Our risk engine, Wings, flagged anomalous behavior in a Firefox extension called Free VPN Forever.*

It should come as no surprise that malicious FreeVPN offerings are beginning to crawl out of the woodwork as the UK, the EU, Australia, and various U.S. States such as Texas and Mississippi begin limiting who can access their services based upon their location. Koi continues, writing:

*The Firefox extension was reading its own logo file, standard behavior, but then doing something unusual with the raw bytes. When we dug into the code, we found a hidden extraction routine. The extension wasn't just displaying the logo. It was searching through the logo's image data, looking for a marker that shouldn't be there.*

*Inside that friendly little planet icon, past where the image data ends, we found malware embedded in the bytes of the PNG image file itself, waiting to be extracted and executed.*

*Free VPN Forever has been on the Firefox Add-ons marketplace since September 2025, with over 16,000 downloads and installations. It's still live as of this writing. And it's not alone, the campaign spans 17 Firefox extensions with over 50,000 combined downloads and installations. Extensions promising free VPNs, translation tools, weather forecasts, ad blocking, the usual lures. What they actually deliver is a multi-stage malware payload that monitors everything you browse, strips away your browser's security protections, and opens a backdoor for remote code execution.*

Okay. Clearly, since PNG images are defined to contain non-executable image data, the authors of this malware must have assumed, apparently correctly, that files of type PNG would not be closely scrutinized by anti-malware scanners and would be allowed to pass. Koi writes:

*When Free VPN Forever loads, it fetches its own logo file logo.png. Standard behavior for any extension. But then something unusual happens.*

The code starts searching through the raw bytes of the image, looking for a marker: three equals signs (===). Nothing after that marker is image data. It's malicious JavaScript, hidden in plain sight. The technique is called steganography - hiding information inside something that looks completely innocent. Security scanners examining the extension's JavaScript files won't find the payload. Code reviewers won't see it. The logo displays normally in your toolbar. Nothing looks wrong. But every time the extension loads, it extracts that hidden code and runs it.

The code pulled from the logo isn't the actual malware. It's a loader, a small program whose only job is to fetch the real payload from a remote server. The loader reaches out to [www.liveupdt\[.\]com](http://www.liveupdt[.]com). If that fails, it tries the backup: [www.dealctr\[.\]com](http://www.dealctr[.]com). The request includes a signature parameter, so the attackers can track which infected extensions are checking in.

Primary: [www.liveupdt\[.\]com/ext/rd.php?f=](http://www.liveupdt[.]com/ext/rd.php?f=)  
Backup: [www.dealctr\[.\]com/ext/load.php?f=svr.png](http://www.dealctr[.]com/ext/load.php?f=svr.png)

But the loader doesn't phone home every time. It waits 48 hours between check-ins. And even then, it only actually fetches the payload 10% of the time. The other 90%? It just... doesn't. Random chance. This is deliberate. Security researchers monitoring network traffic might watch an infected extension for hours and see nothing suspicious. The malware is patient. It knows that inconsistent behavior is harder to catch than consistent behavior.

When the payload does arrive from the C&C server, it's not readable JavaScript. It's been transformed using a custom encoding scheme. The decoding algorithm is almost playful in its simplicity:

- Swap all lowercase letters to uppercase, and vice versa
- Swap all 8s and 9s
- Base64 decode the result

The decoded payload gets XOR encrypted using a key derived from the extension's unique runtime ID, then stored in local browser storage. Persistence achieved. Now it gets interesting.

The final payload pulled from the C&C server, decoded, and executed, is a comprehensive toolkit for monetizing your browser without your knowledge using affiliate link hijacking. The malware watches for visits to major e-commerce platforms. When you click an affiliate link on Taobao or JD.com, the extension intercepts it. The original affiliate, whoever was supposed to earn a commission from your purchase, gets nothing. The malware operators get paid instead.

It's invisible to the user. You still end up on the product page. You still make your purchase. The only difference is who gets the commission. And then there's the Tracking Injection:

The malware injects Google Analytics tracking into every page you visit. Tracking ID: UA-60144933-8. It collects your extension installation date, how many days you've been infected, which merchant networks you visit, and a unique identifier tied to your browser. Hidden HTML <div> elements get injected into pages with IDs like `extwaigglbit` and `extwaiokist`. These elements contain tracking attributes, installation days, signatures, merchant network data, that can be read by scripts on the page or by the extension itself.

You're being profiled, and you'd never know it.

*Then there's the Security Header Stripping: The malware actively removes security headers from HTTP responses:*

- *Content-Security-Policy - gone*
- *X-Frame-Options - gone*

*These headers exist to protect you from clickjacking and cross-site scripting attacks. The extension strips them from every response, on every site you visit. Your browser's security model is quietly dismantled.*

*And then there's the CAPTCHA Bypass: The malware includes multiple methods for bypassing CAPTCHA challenges. One method creates an invisible overlay and simulates user interaction. Another loads an external CAPTCHA solver from [refeuficn.github.io](https://github.com/refeuficn). A third checks if you're logged into Baidu and uses your account status as verification.*

*Why would malware need to bypass CAPTCHAs? Because some of its operations, like the hidden iframe injections, trigger bot detection. The malware needs to be able to prove it's "human" to keep operating.*

*What was that about Hidden Iframe Injection? The extension injects invisible iframes into pages, loading URLs from attacker-controlled servers. These iframes enable ad fraud, click fraud, and additional tracking. They're created, used, and deleted, leaving no visible trace. Referrer policy gets manipulated to hide the traffic's source. The iframes disappear after 15 seconds. Forensic analysis would need to catch them in the act.*

*What makes GhostPoster effective isn't any single technique, it's how they're layered together. Steganography hides the initial payload where scanners won't look. Staged loading means the actual malware never exists as a file, it's fetched at runtime. Custom per-browser encoding defeats pattern matching. Random delays and probability checks make behavior inconsistent and harder to observe. Time delays prevent the malware from activating until 6+ days after installation, long after most security reviews would have concluded. XOR encryption protects stored data from casual inspection. Each layer isn't particularly sophisticated on its own. Combined, they create something genuinely difficult to detect.*

*Free VPN Forever isn't alone. We found 16 other Firefox extensions communicating with the same C&C infrastructure - [liveupdt.com](https://liveupdt.com) and [dealctr.com](https://dealctr.com). Different extensions, different lures, same backend. Some use the PNG steganography technique. Others download JavaScript directly and inject it into every page you visit. Others use hidden `eval()` calls with the C&C domains encoded using custom ciphers. Same attacker. Same servers. Different delivery mechanisms.*

*This looks like experimentation, testing which approach evades detection longest, which gets the most installs, which generates the most revenue. Collectively, these extensions have been installed over 50,000 times.*

*And GhostPoster isn't the first time we've seen free VPN extensions turn malicious. It's becoming a pattern. Earlier this week, we exposed Urban VPN Proxy - a Google-featured extension with 8 million users that was secretly harvesting AI conversations from ChatGPT, Claude, and Gemini and selling them to data brokers. Before that, FreeVPN.One - another featured, verified extension with 100,000+ installs - was silently capturing screenshots of everything users browsed, including bank accounts, private photos, and sensitive documents.*

*Free VPNs promise privacy, but nothing in life comes free. Again and again, they deliver surveillance instead. What makes GhostPoster dangerous isn't any single technique. It's the access. These extensions strip your browser's security headers on **every** site you visit. They inject code into **every** page. They maintain a persistent connection to attacker-controlled servers, waiting for instructions. The payload can be updated at any time. What runs in your browser tomorrow is entirely up to them.*

*The steganography is clever. The layered evasion techniques show operational maturity. But the real threat is simpler: 50,000 users installed extensions that gave attackers full control of their browsers, and those extensions are still live on the Firefox Add-ons marketplace.*

To give everyone an idea of the sort of extensions, they list the names of the actual ones they found:

- free-vpn-forever
- screenshot-saved-easy
- weather-best-forecast
- crxmouse-gesture
- cache-fast-site-loader
- freemp3downloader
- google-translate-right-clicks
- google-traductor-esp
- world-wide-vpn
- dark-reader-for-ff
- translator-gbbd
- i-like-weather
- google-translate-pro-extension
- libretv-watch-free-videos
- ad-stop
- right-click-google-translate

Needless to say, nobody wants this sort of crap lurking inside their browser and tremendously reducing its native security guarantees by removing all incoming website security measures which prevent all manner of hijinks. We've seen that movie. It doesn't end well.

There's not really anything anyone can do. So the original admonishment is still operative and still applies: Don't just rummage around adding every random browser add-on that presents itself. Do everything you can to limit your usage to those that you really need.

