ICANN86 Seville | PF – SSAC Lightning Talks (1 of 2)
Wednesday, June 10, 2026 – 14:45 to 16:00 CEST

KATHY SCHNITT

Hello, and welcome to the first of two SSAC lightning session talks. My name is Kathy, and I'm the participation manager for this session. Please be advised that this session is being recorded and is governed by the ICANN Community Participant Code of Conduct, the ICANN Expected Standards of Behavior, and the ICANN Community Anti-Harassment Policy. Regarding participation today, this session is designed for internal discussion among SSAC members on current security and stability and resiliency topics. Observers are welcome to watch, but this session is not open for observer participation. For SSAC members, all members have been promoted to panelist status in Zoom. To join the speaking queue, use the raise hand feature, even if you are physically present in the room. When called upon, please state your name for the record. You may use the Zoom chat with one another, but please note that it is visible to observers in real time. For observers, the speaking queue is limited to SSAC members, and observer chat has been disabled for this session. I will now hand the floor back over to Barry Leiba.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file but should not be treated as an authoritative record.

BARRY LEIBA

Hi, welcome to the SSAC Lightning Talks. As everybody knows by now, my favorite part of the week. We have three good talks for you. I will take questions from observers if there's time after the SSAC questions, so think about what you want to ask, if you want to ask anything, and we'll see if we have time for you. First talk is from Greg Aaron. How's-- What? You are... Kathy's... She's teeing up my... Okay, we're going to switch the order from what I have. Raffaele, we'll take you first. "Disrupting the Internet in the Name of Copyright: An Italian Story." Go for it, Raffaele.

RAFFAELE SOMMESE

Great. Okay, welcome, everybody. So this is a presentation that I wanted to do because of the work that SSAC has done in regard of DNS blocking and because this was a paper that we wrote, actually last year, about the mess that is happening in Italy with DNS and IP blocking for reason of copyright. But I realized when I was making the slide for this session that this probably is not only an Italian story, given that we are in Spain, and Spain is facing the same problem. But it's a story of how a small group of influential people can convince basically government to do very draconian filtering rules for standard of the Internet as we know nowadays, and how this can easily go wrong, all in the name of football, with similar legislation that have been proposed across Europe. In Spain there is LaLiga, there is France, there is Germany, and apparently there's US coming up with this. But how this platform called Piracy Shield was born.

So in July 2023, the Italian government proposed this law, the 93/2023, and this law was approved unanimously by Chamber and Senate. This is something that rarely happens in Italy, that everyone agrees on something. And basically, the law allows copyright holders to request blocking an IP address, IPv4 address, or domain names involved in the sole activities of illegal football streaming within 30 minutes of detection. And AGCOM, that is the national regulator from Italy, oversees and provides a platform to communicate these blocks. The platform was donated by a third-party society, SP Tech, a branch of a legal counseling organization.

BARRY LEIBA

Go ahead. Keep talking. No. We're trying to get rid of your screens, but... That one's smaller. It's now small enough it's not too annoying. Yeah. Drag it up to out of, to the right. Yeah. Cool.

RAFFAELE SOMMESE

Okay. Hide video panel. Yes. Okay. So initially the platform was supposed to support only 20,000 FQDN and IPs to block because of technical reasons and because it was a burden to the operator. And the operator needs to provide their own blocking platform, so they had to implement, for example, BGP /32 blackholing for the IP address or DNS filtering on their authority name server. The way it works is that copyright holders have the right to insert tickets, and they contain this IP address or FQDN to block, and ISPs query the system, download the list of tickets to block, and apply the filtering, all within 30 minutes. They cannot exceed this threshold. What if

they make a mistake? It's a very short time. Can I complain? You can't. Maybe in very few cases, but there is no formal way.

So some bumps down the road. On 1 February 2024, an IP address from Cloudflare made it to the list, and tens of thousands of websites were unreachable for around 40 hours from Italy. Then Zenlayer and Imperva, two prominent CDNs, also made it to the list, and drum roll, on 19 October at 18:56, drive.usercontent.google.com made it to the list, and you can imagine what happened. Google Drive stopped working in Italy for more than basically a day, because sure, the block was removed right after, at 00:20 of the next day, but DNS has this weird attitude of caching things, so things were bad for an entire day. And you can imagine, "Oh, they implement this nice platform. At least it worked, right? We had an increasing number of legal subscriptions." Well, the answer is no. The number of legal subscriptions stayed exactly the same of 2023.

But then they said, "Okay, let's make a new law. Let's make an amendment to the original law." And now the law says, "You know what? IP addresses should not be used for the sole activities of illegal streaming of football matches. They can be used predominantly." But they didn't define what predominantly means. But they made an official unblock concept, that basically you can request within five days that your address has been blocked or your domain has been blocked to be unblocked if you think that they are wrong. But the list is not public, so how do you know? And now also VPN and DNS providers, recursive resolvers,

regardless of where they are, regardless in which jurisdiction they are, they are subject to apply this filtering. And if an ISP has any suspicion of illegal activities from a user, they need to contact the police. Otherwise, they risk one year in jail. This is in the law.

So you may think this is a success. It's unvetted blocking powers granted to private entities, because copyright owners are private entities, with no clear explanation or review process of the block, with a lack of transparency because no public list is out there, and with a lot of collateral damage, as we saw with Google Drive. And despite all of this, AGCOM, the national regulator, says that Piracy Shield is very effective because it blocks a lot of IPs and a lot of domains. If you think about it, it's like a law is effective because it puts a lot of people in jail. But what lies behind all these blocks?

So the problem is you cannot reconstruct easily the collateral damage because the list of blocks is not public. And to overcome this, what we used is that someone on GitHub leaked the list of blocks. But you cannot publish an academic paper out of someone leaking something on GitHub. So we validated basically the entries because AGCOM gives you a portal where you can insert a single IP or a single FQDN and basically check whether it's blocked or not. And it has CAPTCHA. So you can imagine how this went. Someone was wrong on the Internet, and you spend an entire night to prove that that person was wrong. So we reconstructed the list, and we reconstructed basically the blocking activity from February 2024 to June 2025, showing 10,000 IPv4 and 42,000 FQDNs blocked, originating from almost 4,000 blocking requests, with 98% of IPs

and 44% of FQDNs still blocked as of June 2025. So as you can see, there is not that much unblocking activity going on. And of course, it's football, so blocking activity was during weekends. And the data set does include data before 2024, where the platform was still active, but it's still a solid base to understand what was going on.

The first interesting result that we find out is that, out of all the IPs blocked, 77% were geolocated within the European Union and 38% in the Netherlands. You can literally go after these people because they are within Europe, within the remit of Europol. And a single provider alone accounted for more than 9.5% of all the blocked IPs, concentrated in 15 /24s, so they just migrate from one IP to the other. And OVH, Scaleway, Hetzner, other providers that show up on this list basically were more across different /24s, suggesting basically there was a sort of reuse of shared infrastructure from streamers. So the only provider that seems to have luck in getting unblocked was OVH, because basically there were a couple of unblocked IPs later on from OVH, possibly indicating that basically there was a benign reuse that was noticed and unblocked by the platform. And the unblocking of FQDNs instead was more like, do you remember initially I told you that there was a limit of 20,000 FQDNs? That limit was reached pretty quickly, so they start basically to unblock the old resources, regardless if there was still malicious activity or not on those older resources, because they had a limit in the platform. And the interesting thing is that only 51% of the blocked IPs still respond to probes. So also in the case

of IPs, there is a lot of abandoning of old resources. So what about collateral damage?

So the first thing is that we noticed that a lot of these IPs that were blocked were leased. So we used methodology that some of my colleagues developed to infer leased prefixes, and we found 24% were leased. And the thing is that streamers can easily rent an entire /24, abuse that, and if they get blocked, they migrate to the next one. And 4%, actually, of these leased IPs were released to other companies after the blocking, and 250 IPs were actually reused by legitimate services. So you have out there unsuspecting businesses that are acquiring blocked and unusable resources on the Italian market. The second collateral damage that you have is, of course, we have a lot of shared infrastructure like Cloudflare and all the major CDNs, but also normal providers. There is this wonderful concept of virtual host, and collateral damage may happen at several layers. And you are experts in this matter, so I don't need to go through extensively over this figure, but if you block an IP address, you can block several services of domain names or several services of a company. You can block their name server, their email, their A record, and so on. So we use the data that we collect in OpenINTEL. It's a large-scale measurement platform, and CT logs data, so domain names that we learn from CT logs, to understand, okay, how much collateral damage is happening due to this block of IPs?

We analyzed 262 million domain names and 1.8 billion FQDNs in search of this collateral damage and found 7,000 FQDNs

collaterally damaged by Piracy Shield. Among around 2,000 responded to HTTP and HTTPS requests. For sure, because we manually verified again, 510 of these websites were non-streaming related, so they were definitely collateral damage. 617 of them were streaming related, so there was a good reason for blocking the IP. The interesting part is that only a small fraction of the blocked IPs were responsible for the majority cases of collateral blocking, because you need to block a single large provider to end up with blocking tons of websites. And most of the legitimate affected websites were in languages of the European Union, actually, so French, Spanish, German, and Italian. And one notable case that we noticed during the analysis was that 19 legitimate Albanian websites were hosted on a single IP address of WIIT Cloud, and they were completely unreachable from Italy.

So looking at historical collateral damage, we also saw a similar picture. So 7,000 FQDNs and 665 of which non-streaming related. One very interesting collateral damage was that we found an IP address of a VPS at Hetzner that was rented by a Portuguese hosting provider, disrupting 325 of their domains. 169 of these domains also relied on the IP for email and web hosting. And the company was completely unaware that they were blocked because of Piracy Shield. They thought that was some problem of Hetzner. They were unable to route to Italian customers, and they just noticed that because they were sending invoices to their Italian customers, and they were not getting any response from them. And they had no idea that this was due to Piracy Shield.

So another collateral damage that we found is that we looked at how many of these IPs that they blocked were anycast, because why do you want to do anycast of an illegal streaming platform? And 176 anycast IPs were blocked, and most of these anycast IPs, of course, belong to anti-DDoS platforms like StormWall, DDoS Guard, X4B, that of course faced unintended disruption. The problem of DDoS is that sometimes you adopt a DDoS provider only when you are under attack, so you change the IP only when you are under attack. So the counterintuitive result is that when they were under attack, they became unreachable from Italy. And we found also an interesting case of collateral damage involving an anycast Google IP that apparently was used by Telecom Italia to serve the blocking page for the FQDNs that were blocked by Piracy Shield. So the platform managed to block itself.

The last collateral damage, of course, relies on expired FQDNs because domains expire. And 10% of the 18,000 still-blocked FQDNs, we found, were unresolvable. The majority of the ones that were unresolvable basically tend to have an earlier blocking date. That means that they somehow get abandoned by streamers. And among the 24,000 that get unblocked, 34% were still resolvable. So we don't know again if they get unblocked because they were not hosting illegal activity anymore or just because of the insertion date. And based on RDAP data, 119 FQDNs were re-registered. So there is someone else that now owns these domain names that is not reachable in Italy.

And what streamers did in face of this platform, basically, they found a loophole because the platform, while it supports IPv6 blocking, doesn't block IPv6. And guess what? A lot of them start to serve over IPv6. A lot of them start to serve over a new IP, and basically, they start to evade because IP addresses and domain names are very cheap. You can easily migrate to a new one, and blocking them does make no effort for these illegal streamers to get to a new infrastructure.

So what's the key takeaway of this study? The platform is causing collateral damage still nowadays, and no action has been taken by AGCOM or the Italian government in face of this study. IP-level blocking is indiscriminate and must be avoided. And even if you really want to do FQDN blocking, this should be used as last resort and with a tight time window. The argument for the 30-minute blocking is because you need to prevent that illegal football streaming is watched by someone, and you cannot wait for a judge. But a football stream lasts 90 minutes, not forever like the platform blocks the resource right now. And in general, operators should be informed of the fact that they have been blocked. At least they can complain, or they can report the abuse if it's one of their customers to their national authority. And given that most of these resources are within the European Union for reason of latency, we can do probably a better approach with Europol to try to take down these services. In general, can we fight piracy without harming the Internet? Thanks, and this is the link to the paper if you are curious.

BARRY LEIBA All right. Thanks, Raffaele. We have a few minutes for questions. Are there any?

RAFFAELE SOMMESE Oh, yes. First in the queue, we have Peter Thomassen.

PETER THOMASSEN Hello. Peter Thomassen. So there is this piracy site, it's called, like for movie streaming, it's called the General Television Lending Department, and they run on gTLD-servers.net. And so what happens if I get that domain on the blocklist? To be clear, that's involved in resolving .com and .net domains.

RAFFAELE SOMMESE Again, the problem is the list is not public, so there is not even a way to vet if this kind of action happens and how the...

PETER THOMASSEN Oh, sorry, I think I was too implicit. The gTLD-servers.net domain is a server that receives DNS queries from resolvers when they try to resolve any .com or .net domain. So when I block that, will that mean that under Italian law, all .com and .net domains will be blocked because I claimed that domain runs some streaming platform?

RAFFAELE SOMMESE Yeah, I know. If the operator implemented the law literally, you need to block basically all the requests, even the internal one of the recursive resolver to that domain. You cannot escape from that. The thing is that, of course, operators are rational people, so they try to avoid this kind of thing, and they block the last query that comes from the user. But ideally speaking, yes.

DANIELLE RUTHERFORD Okay. I think we had somebody else, and then Benedict. Warren? Were you in the queue? Yeah, the hands got accidentally lowered. So Warren, and then if we have time, Benedict, but we're running a little short on time. We've got some spillover time for questions after all three presentations.

WARREN KUMARI I'll go fast. So yeah, a friend of mine got a /24 and started up a small company in the US. And he discovered that a bunch of Italian folk couldn't reach him. And it seems like there are seven or eight addresses that are out of the /24 that were previously used by someone else. And I was like, "You could get those fixed." And he was like, "I don't actually care if Italian folk don't reach me." So this seems like it's not just collateral damage to that, but it's also affecting some set of Italian people. But at least you've solved all of the streaming things, right? Mission accomplished.

RAFFAELE SOMMESE Some people may not care. But if you are operating a commercial service, especially if you are in the EU and most of these IPs are in the EU, you may risk really damage in terms of economical damage, because again, if you don't send invoices, you don't get money.

WARREN KUMARI Oh, I fully agree. His service, I think, would be of benefit to Italian people, but they don't have access to it. And he's like, "Well, they shot themselves in the foot."

RAFFAELE SOMMESE Yeah.

PETER THOMASSEN Hmm.

DANIELLE RUTHERFORD All right. We are actually at time, so Benedict, please hold your question until the final 10 minutes of the presentation. We'll come back to that, and I'll keep track of that.

BARRY LEIBA Okay.

DANIELLE RUTHERFORD Barry.

BARRY LEIBA Okay, Chaoyi, you're next. We have "DNS Security Academic Papers Published in 2025." More interesting than you might think from the title. Go for it.

CHAOYI LU All right. So good afternoon. So it's been another year. I did this last year, and I remember clearly I was asked to do this again. So here it comes, summarizing DNS and security academic conference papers published in 2025. So I limit the scope a little bit to academic conference papers only. So we'll look at what has been focused by the academia, what has been found about the DNS in the last year, and how they could interest our community. So same as last year's presentation, our scope includes 12 Tier 1 and Tier 2 security and networking venues. So counts and numbers to start with. This year, our investigation found 20 papers about DNS and its security. So recall the number in 2024 was 23. So we can conclude that DNS is still a hot research topic in the academia. The papers now show an even distribution among three categories, measurement, new threats, and detection or solution. So here are the papers in the measurement category. They focus on obtaining real-world data about different types of DNS abuse, DNSSEC errors and new features, and servicing and using DNS data for other types of measurements. In the category that discovers new security threats or attack vectors, denial of service is still a popular subcategory, followed by cache poisoning and registry operational flaws. And

the last category, papers proposing novel mechanisms for detecting DNS abuse and introducing infrastructure solutions leveraging DNS features.

So next, I will try to walk you through the core technical insights and major conclusions from these studies. And again, full disclaimer, I did not write any of these. So the contents presented here are taken from the papers and are of the authors' opinion. So let's start with the measurement papers. Well, this paper tries to find key factors behind phishing domain registrations. So we know that cybercriminals are continuously registering new domains that are highly concentrated in a handful of registrars and TLDs. So supported by ICANN's INFERMAL project, the authors quantify how operational and financial features are driving this abuse. Excuse me. So conclusion number one, price matters. They find that lowering domain prices by $1 correlates with a 49% increase in malicious registrations. And also bundling free DNS services causes a massive surge. Conclusion number two, automation is heavily exploited. Registrars that allow bulk automated registration via APIs experience a staggering increase in abuse. And conclusion number three, verification is effective. So even enforcing simple email or mobile verification drops abuse by 70%.

Okay, this paper investigates the lifecycle and infrastructure of APT domains. Well, detecting APT domains has been difficult because the relationship between malicious infrastructure and their actors is highly transitory, and attacks can persist for years with dynamic IP changes. So by analyzing the DNS traces of hundreds of known

APT groups, the authors find that APT actors are frequently reutilizing their infrastructure. Well, they also highlight the importance of historical data in forensics because most IPs will stop pointing to active APT domains after the campaign is disclosed. So for organizations, the authors find that APT attackers typically provision their domain infrastructure more than 300 days before an attack is reported. So they arrive that the historical network logs are recommended to be maintained for at least two years to effectively detect and trace these threats.

And next we look at how enterprises are defending themselves from DNS abuse. Well, this paper studies the defensive registration practices of Fortune 500 companies. So enterprises can proactively register variations of their brand domains, like in the forms of TLD squatting, typosquatting, bitsquatting, or homophones to protect their trademarks. But this study finds that the overall level of enterprise engagement in protective registrations is still very low. Over 200 major companies registered only fewer than 10 protective domains each, and most of them are focusing on one subcategory, which is typosquatting. And meanwhile, the current online brand protection tools, or OBP, are very highly used, often miss critical domains that still receive high user traffic. So the authors strongly recommend that companies utilize passive DNS data to identify and proactively register available domain variations that see high query volumes.

And moving on to measuring DNSSEC. This paper analyzed DNSSEC misconfigurations at scale using years of DNSViz logs. So given that

DNSSEC misconfigurations can make domains unresolvable in the worst case, which are the frequent errors, and how can we fix them? So the authors find that when breaking down the errors, NSEC3-related errors are the most prominent in DNSViz data, particularly due to non-zero NSEC3 iteration counts or missing proofs of non-existence. And following that, delegation and signature errors, such as invalid KSK algorithms in DS records or completely missing and expired signatures, also contribute. And so based on these, the authors also built an open-source framework called DFixer that automatically generates commands for fixing these prominent errors they discover.

Another paper measured the adoption of automated mechanisms like CDS, CDNSKEY resource records, and authenticated bootstrapping, or AB, which are designed to automate DS record updates and initial key configurations. So after scanning millions of zones, the authors discovered that only 5.5% of scanned zones are correctly signed. But interestingly, CDS deployment is gaining primary traction among smaller, localized DNS operators, particularly in Switzerland, which is likely driven by the original financial incentives. And while authenticated bootstrapping is still not very widely adopted, it is implemented with near-perfect correctness whenever it is used. So pointing out that the low baseline of DNSSEC deployment itself remains the primary bottleneck.

And finally, in the measurement category, this paper investigates Internet readiness for DNS over IPv6. So RFC 3901 has historically

made IPv6 support optional for the DNS, and current data shows that less than 60% of all DNS zones are supporting IPv6-only resolution. So is DNS impaired by IPv6, and is it time to require IPv6 for the DNS? So by tracking the daily resolution of the top 10 million domains under various path MTU discovery scenarios, this study finds that while DNS timeouts are slightly higher in IPv6 due to some local network misconfigurations, IPv6 and IPv4 fragmentation are comparably functional. The traditional fear that broken MTU paths would severely impact IPv6 DNS resolution is not indicated by their findings. And so the authors conclude with an active policy recommendation: it is time to officially recommend that IPv6 should be used in the DNS.

So our second category focuses on new DNS security threats, and this year it is mainly about implementation issues. So we start with a denial-of-service vector that, I quote from the authors, "turns a security feature into a weapon." So in DNSSEC, the checking disabled, or CD bit, is designed for troubleshooting. When this bit is set, a validating resolver returns unvalidated data and signatures without authentication. However, the authors discovered that many resolvers are caching and actually reusing this unvalidated data, returned by CD-bit set queries for subsequent DNS resolutions. So with CD bit set, an attacker can send a query to a resolver and then inject a forged invalid DNS response. The resolver then caches this bad data. And then when an ordinary client sends a standard query for that victim domain, the resolver will reuse the unvalidated cache data, fails validation, and returns a SERVFAIL

error, cutting off access to the victim domain. So the authors find this threat widespread, affecting major DNS software packages and public DNS services, highlighting an urgent need to restrict the reuse of cached, especially unvalidated data.

Another denial-of-service threat comes from devices called transparent DNS forwarders. So these are devices, mostly consumer routers, that forward DNS queries upstream without rewriting the source IP and without caching the responses. Because they do not respond directly, they are invisible to standard scanning platforms like Censys. But studies have found that they do account for 30% of the open DNS infrastructure. So by scanning, the authors first find 530,000 transparent forwarders and demonstrated how they can be weaponized. Specifically, they can be used to penetrate network firewalls to reach what is called shielded resolvers behind a firewall. Because these internal resolvers are hidden, they often have much looser rate limits or even support massive response types like ANY queries, so allowing attackers to generate highly amplified flood traffic far more efficiently.

The next paper introduces a flush-reload side-channel attack targeting DNS forwarders, using their caches to precisely infer whether a user has visited a specific website. So the attack works in three steps. First, the attacker flushes and clears the forwarder's cache, and second, they wait for a specific time interval, and third, they will measure the response time for a target domain query. If the response was extra fast, it means a local user queried it during

that interval, and there is cached data in the forwarder. So the authors tested this across various browser and OS combinations, showing that some are susceptible to this cache-snooping technique.

DANIELLE RUTHERFORD Five minutes.

CHAOYI LU Ah, okay. Also about caching, this paper demonstrated that a classic threat is back: DNS cache poisoning via the birthday paradox. So the gist is that there is an implementation blind spot after ECS is introduced. So specifically in these implementations, queries with ECS are aggregated based on a triplet of QNAME, QTYPE, and subnet. So an attacker can use this to bypass query aggregation by sending queries with different subnets. So according to the birthday paradox, the success rate to match the transaction ID hits over 99% after just nearly 2,000 rounds.

And to conclude the threats category, this paper explored vulnerabilities at the TLD registry operations. So while domain lifecycles are standardized via the EPP protocol, individual registries differ in their practical implementations. So the authors uncovered major flaws that allow for permanent free domain takeovers. For example, in internationalized domain names, registries often automatically create different script variants that relate to each other, like this example. We have traditional and

simplified Chinese characters together, so when one of them is registered, the other will be automatically registered by some registry, but they often fail to properly sync their lifecycles. And also, we see some of the glue records, 23% of analyzed glue records have been completely abandoned. So attackers can permanently control these relic domains without paying, providing a completely free infrastructure for illicit activities.

And the final category, some detection and solution works. This paper introduces a system called MANTIS, designed to catch zero-day malicious domains at the very moment of hosting deployment before traditional methods that rely on user infections or TLS certificates. So it leverages a key observation that attackers consistently reuse low-repute malicious hosting infrastructure. So it uses a GNN encoder. It classifies malicious domains with high accuracy, outputting a reliable daily blocklist that catches threats before users even access them.

And addressing the resurgence of cache poisoning, this paper proposed POPS, or Cache Poisoning Prevention System, designed as a lightweight intrusion detection module for resolvers. So what the authors do is first systematically categorize decades of cache poisoning attacks into several categories and map out their use of port guessing, fragmentation, or out-of-bailiwick data injection. And accordingly, the system enforces three simple but powerful behavioral rules. Rule number one: alert if the number of incoming responses differing by only one single field passes a strict threshold. Rule number two: alert immediately upon seeing the

first fragmented packet of any DNS response. And rule number three: alert if the query name falls entirely outside the bailiwick of the responding server. So once an alert is triggered, the mitigation is straightforward. The resolver forces the connection to fall back to TCP by setting the truncation bit.

And finally, let's look at how we can leverage the DNS infrastructure to solve broader Internet problems. This paper tackled the scalability problem of the Online Certificate Status Protocol, or OCSP, for TLS certificate revocation. So currently, CAs delegate their OCSP hosting to external CDNs, which becomes fragile centralized services. So what the authors proposed is called RevDNS, which is a decentralized CDN-independent revocation scheme built on top of DNS. So a CA will now publish its revoked serial numbers directly in DNSSEC-signed TXT records on their own authoritative servers. So the design utilizes NSEC aggressive negative caching for more efficiency. So if a certificate is valid and not revoked, the resolver can prove its non-existence using NSEC records without contacting the CA. And their evaluation shows that for massive CAs like Let's Encrypt or GoDaddy, storing revoked certificates via ECDSA signatures in the DNS adds a manageable size to zone files. And also, due to DNS caching efficiency, local resolvers can answer over 99% of revocation queries directly and remove the operational dependency on third-party CDNs. Okay, so that's a wrap of the past year's academia about the DNS. Thank you.

BARRY LEIBA Perfect timing. So we don't have any time for questions on this right now, but save them for the end bit. I'll just add that with NDSS, after many years of it being the last week in February in San Diego, next year's will be in Seoul at the end of March. So now Greg Aaron will tell us how cybercriminals buy domain names.

CHAOYI LU Oh.

GREG AARON Okay. Hello, everyone. This is a paper that we wrote through Interisle Consulting. One of my co-authors, Karen Rose, I think is behind me. And we wanted to look at what's going on in the world. This community and others need information. That's a good way to make decisions about things. So we looked at four questions. We wanted to look at domain names registered in 2025. These are the new creates, the names that are being created and sold. We're not looking at older, renewed domains. But we wanted to see what happened with those domains and maybe learn some things about the domain abuse problem, put it into context, put some numbers to it. And we wanted to know what percent of those registrations got involved in abuse. Where were they concentrated? We'll talk about associated domain checks, and then what's going on? What's happening and why?

So to do this, one of the standard ways of looking at abuse and counting it is to look at reputation blocklists. This is a standard technique that's used by academics. It's also used by ICANN's OCTO team, for example, to run the Metrica system. These are domains that are getting flagged for abuse, and they are a real-life representation of what's being rejected out there by people who are using the Internet, basically. And of course, there probably isn't a better source, because there aren't companies who have a view of what's going on in the world who share data. But we put together a bunch of blocklists, and we have overlap with some of the other folks who do this kind of measurement. All the data is available. It can be obtained. So in the back, if you ever want to get into the methodology, talk to me, and there's a full explanation in the back.

So what percent of new registrations were involved in abuse, supposedly sold to bad actors? There were 85 million new registrations in the gTLDs in 2025. Those are the new domains that were created. As of the end of last month, 8.5 million of those had appeared on the blocklists we monitor. So that's 10% right there. That's the floor. The domain names registered in the second half of last year can still get blocklisted this year as they reach the end of their first year. Based upon what has happened in previous years, we expect the 10% to rise by 12% by the end of this year. Then we have to think about, well, what do the blocklists miss? They are having a certain view into the Internet. They don't see everything that happens. They're not listing all of the domains that are probably associated with these actors and these activities. ICANN's

security research team in the Office of the CTO did a very interesting study and released it a few months ago, and they did associated domain checks. They said, "We're going to look at a quarter. We're going to see what got blocklisted, and then we're going to see what other domains were associated with those." And basically, what they found is for every three domains they found on the blocklist, they found another two. They also made this a very conservative estimate, they said. For example, they didn't count batches over 1,000 domains. And for those of you who have gotten into this world, you know that batches of 1,000 domains actually happen quite a bit. So if you use that ruler, that's going to project to 20% of the domains. And we did some case studies I'll tell you about in a minute that kind of triangulate on that number. So that's the estimate, one in five. That is a large percentage.

And volume of abuse has been increasing. This is phishing specifically. Every quarter in the sources we would monitor, we would see about 100,000 domains reported, but that's increased by eight to nine or tenfold. It goes up and down, but the number of domains getting used is increasing because of automation, and there's a cat and mouse game going on. We try to take down domains faster. The criminals respond with more, and there's more automation. So where did the abuse concentrate? It was widespread in that it appeared in pretty much every gTLD and at almost all of the registrars. There are about 3,000 accredited registrars. It appeared in the great majority of them, but it's also really concentrated in a few spots. So we looked at where did it

appear in terms of TLDs, registrars, and families, because some TLDs and some registrars are actually just controlled by one decision-maker, one company. Here's the distribution. So of all the domain names in the gTLDs, 20% of them are currently registered and existing in the new TLDs. Last year, 42% of the new domains registered were in the new TLDs. So a lot of the activity in the market happened in the new TLDs. 61% of the domains that got blocklisted last year were in the new TLDs. So there's a definite disproportion in the sectors. The new TLDs are often sold, for instance, at lower prices. They are incurring more abuse, and I'll get to that more in a bit.

As far as concentration, there were eight registrars that had over half of their new registrations blocklisted last year. NiceNIC had 88% of its domains blocklisted. A lot of people are not trusting the domains there, or they're showing up in the detection. So that's just blocklisted. There may be more associated domains in their portfolios. So that's not a happy situation. There were five registrar families. So we figure out who owns which accreditations. You can have more than one accreditation, of course. There were five registrar families that accounted for a little more than half of all the gTLD domains blocklisted, Namecheap coming in at number one. They had basically 14% of all malicious registrations created in the gTLDs last year under their Namecheap and Spaceship brands. There were 13 TLDs that had more than 50% of their registrations blocklisted. They're pretty much all new TLDs, some of them smaller than others. And we'll see that some of these TLDs grew a

bit in 2025. They attracted some new registrations, but some malicious registrations evidently came with them. And if you look at the registry families, Verisign comes in first because .com is there. .com is big. Criminals still like .com for some reason, even though it's not always the cheapest of the TLDs on sale, but that's why that chunk is there. .top all by itself is almost another quarter. That's just one TLD, and it has a large number of domains in it.

So what can associated domain checks tell us? We wanted to see what was missed. So when you do an associated domain name check, you want to correlate several features because we don't have contact information. That's an important one. But we can look for associated domains that share the following characteristics. You start with blocklisted domains as your seed. Those are your pointers, where you want to start looking. And then you look for the other domains that were registered at the same registrar within a small period of time. We're usually using registered within 10 minutes, using the same hosting or name servers or A record. And then usually you're going to see that the domain strings follow a pattern. So once you start looking at these, the pattern is often quite clear. So we're looking for some lexical similarities.

So this is the .loan TLD and the dark blue line on the chart. The chart starts in January '24. And .loan was a small TLD and had about 7,000 domains in it for years, and it just was kind of constant and cruised along at about 7,000. And then in August '24, the number of registrations starts to go up sharply, and you see that blue line

slope up real fast. The red line is the cumulative number of blocklisted domains that occurred in the TLD, and it starts to ramp up right along with the growth and is a significant portion of the domains under management. So the TLD grew from 7,000 to over 120,000 domains, and the green line is the renewal rate. The scale's over on the right. So this TLD had a pretty good renewal rate for years, 50% to 70%. That's pretty average for TLDs. And then exactly one year after the growth and the abuse starts, the renewal rate craters down to just a few percent. Basically, we're seeing nobody wanted to renew those domains. It turns out the domain names had been used for criminal activity, and of course, they didn't renew them after they were done with them. Specifically, these domains were registered by an organization known as FUNNULL. That organization runs the criminal infrastructure for a lot of the online criminal activity in Asia. They run the infrastructure for the pig butchering farms. They support cryptocurrency scams. They spread malware and were getting into toolkits to spread malware, and they were involved in sending money to North Korea. These are the worst folks in the world, almost. And they got about 100,000 domains in the TLD. They were the major registrant.

They also got sanctioned in the middle of last year. The US government said, "We're going to put them on our sanctions list, so that means that no one can do business with them in the United States, no US entity." And that was partly to make sure that banks did not do any transactions or processing for this group of people. At the same time, the FBI released a list of domain names that they

had observed on FUNNULL's infrastructure. There are 320,000 of them, and we used those as seed material to look for associated domains. So the sanctions happened, but FUNNULL was still getting domain names after that. They were still registering. In some cases, they were still registering at US registrars. So that was not a cool situation.

Other case studies, we did .pink, which was a small TLD, and it grew a lot, kind of what happened with .loan. We found a 38% expansion rate. Some of those domains were also registered by FUNNULL. The renewal rate dropped from around 65% to 3.1%. We chose .gift because it had never historically had abuse. We wanted to look at a smaller TLD and just see what was going on and one that was kind of quiet from an abuse point of view. It did have some phishing in it, as it turned out. We found a 63% expansion rate from those domains that we had. So that equaled 15% of the registry. It was a smallish, quiet registry. It had a good renewal rate of 61%, so there are registrants who wanted to renew their domains, that was good, but still 15% in that TLD. And .bond had a lot of unusual activity in it, phishing and malware.

Of course, sometimes when you look for associated domains, sometimes it's quite obvious what's going on. This is a good example. Those are the actual phishing URLs, and you see the domains in red. They were advertised through text messages, like the one you see on the right, and they follow that obvious lexical pattern. They were all registered in batches. In this batch, we saw something like 40 or 50 of those domains that got blocklisted. What

did that lead us to? That led us to something like 400 domains. They just go on and on and on. But they were obviously all in the same batch, and we had a high miss rate on the blocklists.

So, the lessons. These are sometimes obvious once you start to look for them, and you can hone in on them. The batches are often very large. They go from hundreds sometimes to hundreds of thousands of domains. The detection and mitigation was partial, probably far from comprehensive. We looked at a couple of points at how many domains that got blocklisted were removed from the zone, so that would've been a registrar or registry hold. The rates range from 7.4% to 16.3%. So not a lot of them ended up getting pulled from the zone. They may have been mitigated perhaps in other ways, but not through EPP server or client hold. What was interesting, though, is when we looked at associated domains, none of them got suspended in these case studies. That could mean several things. Maybe some of them didn't get reported at all to people. But what we see is they weren't getting found somehow and therefore didn't get taken care of.

What we also seem to see is that the payments are probably getting through. So if a criminal is paying for a domain name, one of the things that happens is you get credit card processing fraud checks. Evidently, those are not finding these kinds of registrations because they sit there for a while, and they are getting used. Very few domains are getting suspended in the five-day add grace period, where you might be able to catch some of that fraud through credit card fraud checks. This suggests several things, like

the criminals are paying through debit or gift cards or crypto. These are methods where the payment's going to go through. And I think Tucows said very few, like 2%, of transactions for malicious domains get caught through the fraud checks. So that tells us something about how the criminals are probably paying for this, and that's a really interesting subject for somebody to look into someday.

So what are the economic issues? Why does this happen? How does it happen? It happens a lot. There's a lot of churn. There's a lot of domains being registered. So it's making sense for somebody to take these registrations. The margins are very small, but somebody is willing to accept these registrations on the registrar side. The criminals do represent a reliable source of demand, unfortunately. If it was a big enough problem, we might see more deterrence, but I don't think we're seeing a lot of that right now. So there are places where these sales strategies and volume discounts might even encourage and attract this kind of a registrant. Registries offer rebates and discounts. Registrars often offer those as well. Some of these domains are getting registered at $0.49 retail, and that's where the criminals go. So tolerating that kind of registrant seems to make economic sense, or it's not worth enough time and energy to keep them away. But economically, there's something happening where the incentives are not helping with the deterrence.

Now, I do want to note, there are some registries and registrars who seem to be managing this stuff pretty well. Like XYZ registered

seven million domain names in .xyz last year, which is a large number. That was a significant portion of new sales in the entire market last year. They had a low renewal rate, something like 16%, we calculated. But the number of domain names that got blocklisted was also pretty low considering that high number of registrations. It was about 275,000. It's not a negligible number, but compared to the number of creates they did, it's quite smallish. So something they did over there allowed them to grow, attract a lot of new registrations, but not attract a ton of abuse. And there are other registrars, like Tucows has a very low rate, so you can look at those rates in the paper. Some business models and some management of these businesses is working. So it's not inevitable that the abuse happens.

But what I think we have at this point is we've got so much of this activity going on that it's starting to impose what we call a negative externality. That's an economic term, which is that the criminals are doing great. Somebody's accepting their registrations, but the costs of that are passed on to other people, specifically the public, people who are ending up getting victimized by the criminals. So that's not a good or balanced situation. So the question for this community and others is, what do you do about it? We did not make recommendations in the paper. We wanted to present the information about what's going on. But clearly, prevention is something we have to pay more attention to. Mitigation's good. Associated domain check could be very interesting and have good results. But the increase in sales and also the increase in TLDs,

we're going to have some new greenfields, we're going to have some operators who don't know how to operate TLDs. They're going to do it for the first time. We're going to have some continued problems unless we figure out a better approach to this problem.

So in conclusion, we've got a significant portion of what's going on in the industry associated, unfortunately, with criminal or malicious activity. It's, at this point, I think, an unbalanced situation. It is creating a lot of harm, and we need to stress prevention, measurably reducing the harm and the abuse while letting people get their domain names that they need to, to do the things that are good and beneficial and letting the space and the DNS flourish as we hope it to do. So that is it.

BARRY LEIBA All right. Thanks, Greg. We're going to take questions for Greg's talk for the next couple of minutes, I guess probably just one question, and then we'll go to general questions for any of the presentations.

DANIELLE RUTHERFORD All right, Martin, go ahead.

MARTIN HARRISON This is Martin Harrison, for the record. Hi, Greg. So for .loan, you had a nice slide showing a graph of renewals, and they had a lot of growth. And so to me, it was unclear from the growth rate, or from the renewal rate, how this affected their original portfolio. Because if you have this much growth, then yeah, that distorts the number.

So I guess my question is for .loan and maybe other cases that you studied, how does an episode like this affect their business viability long term?

GREG AARON

Right. So in .loan, I think at the end of that year, after the abuse happened, the only domains that were renewing were the domains that had been in the registry for a long time. Those original 7,000, basically. The normal view, I think, in a lot of registry operators is I like domain names that renew year after year because I have a repeatable and somewhat predictable business. Right? .org is at 80% renewal, .com's at 74%. And so what we're seeing is there's also a different business model where some TLDs are getting along with registrations that only last a year, which is, I don't know if we saw too much of that in the past. I wouldn't be able to quantify it. But normally, if you have a good business, you want to offer a sale special, you want a certain percentage of those to renew at the end of the first year. That was always kind of the calculus. If I offer a lower price, maybe I'll take a hit right now, but then it'll renew at a higher normal price a year from now. And so if maybe 30% of those renew, I'll be in a good spot, and it'll grow my base. So I think that's kind of the economics of what's going on. And some of these TLDs just grew by a lot. Somebody somewhere maybe made some money along the chain, but then everything collapses a year later. Then the question is, do I pursue more bad business?

DANIELLE RUTHERFORD

All right. So it's now into the general Q&A period, so I'm going to go first to the spillover queue, and I've got Jeff, I've got you and Nabil in the spillover queue for your questions for Greg. First was Benedict, but I don't see him on the call anymore. So Warren, I will go to you for your question to Chaoyi. All right. Well, then Jeff, over to you.

JEFFREY BEDSER

Sure. So Greg spoke to economics, and I think we need to take a step back on this particular topic and realize that cybercrime has been measured recently as the third largest GDP after the US and China. There are trillions of dollars. The supply and demand for criminals to get their hands on domains to do cybercrime is not going to go away from any measure. So while it is a big problem, I'm not trying to diminish the size of the problem. I do believe that the most important measurements are going to be around detection, which is not currently required under the obligations. The obligations clearly state a well-evidenced report must be mitigated. It doesn't say anything about detection, though I'm assuming that the ADCs will actually have detection involved, and that it's time to live that's most important, not volume created. I think we need to work toward a world where fast detection results in fast mitigation so that victimization can occur in the window between the delegation of the domain and before it's taken out of practice, whether that be phishing, malware, or what have you.

So I think that one other factor that I did see in the presentation was about the ADCs against the blocklists. And one of the things as a practitioner, working with quite a few parties, I can say now is that most of the blocklists have a domain and an allegation. And most of the parties in the industry have a hard time gathering the evidence to act from that little bit of information. And I would assume that because of the mitigation rates being very low against those blocklists, it's an indication they're not finding the evidence to take action against them because the requirements are a well-evidenced report must be mitigated, not a very loud allegation from a blocklist provider, in my opinion. So, sorry, not really a question, Greg. Just a statement.

GREG AARON

Okay. I think if the problem is reduced to mitigation, that's a losing proposition. Time to live is important, but like with phishing, we see that they're launching more domains, more attacks, because they know that those attacks are going to be up for a shorter period of time. So they advertise them. They may only be up for an hour, but they'll capture maybe a certain number of victims in that time. So that's the cat and mouse game. If you use one domain and get 20 victims, and then you reduce the time to live, and you can launch 20 domains and attacks and get one victim each, you have the same result, and the criminal's happy and he's got what he wanted. And his yield is the same as before. So I think prevention's still important. And as a note, no one is asking registrars or registry operators to take down domains based on blocklists. Now, some of

these blocklists do confirm exactly what's happening. They've got proof of what's happening, but it's not their business model or their role to report things to registries or registrars. So that's one of the holes, right? That's not their job. They're interested in getting their stuff out onto their list so they can be blocked in firewalls and all the other places they're used. So that is a gap, but it's there by the model. Thanks.

PETER THOMASSEN: Just a quick two-finger on that. The yield may be the same, but the cost is 20 times.

GREG AARON: I don't think so, so much anymore. Now, the cost per domain might be 20 times, but the cost of managing the domains is, because of automation and so forth, has gotten really low and almost trivial. So you've got platforms and kits and stuff, and it's much easier to manage a large number of attacks now than it used to be. So the cost of domains, absolutely, I agree. But the cost of everything else has gone down.

BARRY LEIBA: We have about four more minutes for Q&A, so keep your questions and answers brief, and Danielle's running the queue.

DANIELLE RUTHERFORD: Next, we've got Nabil.

NABIL BENAMAR Thank you, Danielle. My question to Raffaele. So in your presentation, I was wondering if you take into consideration the case where some ISPs are providing to their customers private IPv4 addresses. So when the address is blocked, it's the public address of the ISP, and the entire network of this ISP will be blocked. So is it part of the case study?

RAFFAELE SOMMESE You mean if there is an IP that the ISP is using, for example, for outgoing traffic of their customer and that IP gets blocked?

NABIL BENAMAR They call it carrier-grade network.

RAFFAELE SOMMESE Yeah, it's carrier-grade NAT. It's rare because, of course, most of the block happens because of infrastructure that is hosted by somewhere. It may happen, but I don't think it's really hard. And also, most of the IPs are owned by the operators themselves, so they are not reused or things like that. But sure, if you are a small operator and you lease addresses, this is something that can happen.

DANIELLE RUTHERFORD Gabe, go ahead.

GABRIEL ANDREWS

I just wanted to add a bit of flavor to the FUNNULL conversation of those bad guys. And forgive me if this is already well known to those in the SSAC. But one of the things I remember from talking to some of our case agents that were looking at this and that informed some of the conversations we had on that here in GAC and ccTLD operator conversations was that we saw the FUNNULL guys using a technique that I had to be taught on, CNAME forwarding, which sort of starts with whatever domain the victim sees. And so they'll do this long, drawn-out phish over texting with someone for weeks on end before sending them to a cryptocurrency investment site. And so the link to that won't directly go there, but rather it'll hop through multiple layers then of additional domains using CNAME forwarding. But the way that it works, though, is that they'll still only see the original domain in their browser because of how that resolves. And I don't know how well that's understood, but I mention it for those with abuse teams and the registrar colleagues in the room because that challenged some of our investigators. It was obviously an obfuscation method designed to make it harder to see where that domain eventually resolves. But you might have then four domains in the chain before it ultimately resolves to the cryptocurrency investment site. I don't know how that's reflected in abuse reporting, if that's captured, or if it's even observed in your investigations, but I just thought it was worth mentioning. Over.

GREG AARON: Yeah, they were very sophisticated, and they used the CNAMEs and redirection a lot. They also hopped around and mixed up the registrations, which makes it actually harder to do associated domain checks. I describe in the report how they would get a sequence of numbers in a TLD, 0005, 0006, and they'd do 10 or 20 of those in one TLD. Then they'd go do the next sequence in another TLD. Then they'd come back and do the sequence before and really mix things up. Now, if you started to look across TLDs, and it's easier to do this in retrospect, you see what they were doing. But at the time, it was probably pretty confusing.

BARRY LEIBA: Okay. We're out of time. It's time for everybody to go out and enjoy the absolutely amazing tapas that are being served in the... I said they were amazing. I didn't say they were good.

JEFFREY BEDSER: Thank you, Barry.

BARRY LEIBA: Thank you all for coming. Thank you to the observers as well.

[END OF TRANSCRIPTION]