ICANN85 Mumbai | CF – GNSO: CPH DNS Abuse Comm Update - DGAs
Wednesday, March 11, 2026 – 13:15 to 14:30 IST

NIKKI SCHULER

Hello and welcome to the CPH DNS Abuse Community Update. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior, the ICANN Community Anti-Harassment Policy, and the ICANN Community Participant Code of Conduct concerning Statements of Interest.

If you would like to speak during this session, please raise your hand and Zoom. When called upon, virtual participants will be given permission to unmute. On-site participants will use a physical microphone to speak. Please state your name for the record and speak at a reasonable pace. I will now hand the floor over to Dennis Tan.

DENNIS TAN

Thank you, Nikki. This is Denny Stan with the Registry Stakeholder Group. I'm one of the co-chairs of the DNS Abuse Working Group. And sorry, I couldn't be with you in the room, but my co-host, Reg Levy from the Registrar DNS Abuse Working Group is with me. And both of us will try to moderate this session today. So, again, good morning, afternoon wherever you are in the room or remotely. We have a mainly two-part agenda today.

First, we're going to be discussing DGAs, and I'll get in a minute about what we're talking about DGAs, what's the current existing operational framework, and what might be potential improvements to that operational framework. The second part will have Rowena Shoo from the NatBeacon Institute and she will give us a little bit update on how the institute is measuring DNS abuse mitigation.

Before I get into the substance of this presentation, I'll remind you to log into the Zoom room. That's how we are going to be managing the queue if you have any questions throughout the presentation. So, with that, let's dive right into it. Please go ahead. So, DGAs, Domain Generation Algorithms. Next slide, please.

So, before we get into the conversation about DGAs, just to set the stage why we are talking about DGAs today now. So in December 2025, the GNSO Council published their issue report on PDPs for DNS abuse mitigation. In that issue report, they identified a number of issues or gaps for potential work policy or otherwise. All those numbers of issues, three were identified as priority items. One being Associated Domain Check. Number two, Ungated API Access. And number three, Better Coordination on DGA Threat Mitigation.

The first two, the GNSO Council proceeded to approve the initiation of two distinct PDPs, PDP 1 and 2. In fact, PDP 1 just got kicked off during this ICANN85. But now into the DGA conversation. That's one of the priorities for policy work. We being the CPH DNS Abuse

Working Group, we have advocated for potential improvements to the current state of work in line with the narrowly targeted scope improvements that we can do in order to help mitigate DNS abuse or security threats.

So, let's go into the next slide, please. So, what I want to do and what we've experienced through all our presentations dating back ICANN83, there's a lot of sometimes maybe misunderstanding what a DGA is, what is not. So, I want to give this opportunity to take a step back and try to level set the understanding in a non-technical manner so that we can approach this to all audiences that are watching now today live or potentially in the future through the recording.

Okay, so what is a DGA, a Domain Generation Algorithm, right? The name, it implies it's a software program that creates, generates a number of domain names, but not only a list of domain names, right? Typically, they will be accompanied with other attributes such as a date and time of activation. And so, you can think of a DGA and how it is used. It's a tool, it's not the threat itself, it's how they use it.

And so, you can think of, you know, and I'm using here the analogy of the criminal boss and their associates in which this henchman and the criminal boss have a secret rule book and by which they can identify the meeting locations and the times where they can connect to those meeting locations and obtain more instructions

on how they are going to do their nefarious activities and whatnot, right?

In today's world, DGA, they can create thousands, millions of those meeting locations, date, times, and the criminals use these tools in order to evade detection. Imagine not knowing when and where these communications might happen. And for the defenders, it's even more difficult in order to guess where these communications or rally points are going to be conducted.

So right off the bat, there is a symmetry problem here. It's very easy for the criminals to create these random communication points in the thousands, whereas in the defensive side of the equation, it's very difficult to keep track of these changes. So, it's an ever-moving target in order to try to defend them. But when the defenders are able to unlock or reverse engineering, those secret rulebooks, then that's where we can start mitigating this threat in an ad scale. So, let's go to the next slide, please.

So this is just an illustration of how a response playbook works. And here, I'm trying to display a four-stage playbook in which it starts with the identification of the DGA, and that's through analysis of network traffic, but most importantly to reverse engineering the DGA so that we know what exactly the domain names are going to be and most importantly when they are going to be activated for as a communication point. Once we know that information, we can go to the next stage, which is the disruption.

And this is where registry operators such as Verisign, for example, and others can start helping addressing this security threat, which could be malware trying to communicate with the command and control. And we can disrupt the communication in different ways, and I think some of the sessions in this week try to get into that conversation. One typical way is to reserve the domain names.

So, registries, depending on their operational approaches, they can have the facility to reserve domain names such that that domain name is blocked from registration, from the criminal, and they cannot use it. Other ways are to preemptively re-register those domain names again in order to block the registration from the criminal.

Other ways in order that we can disrupt the resolution is by allowing the registration to go through which allows to capture certain digital footprints from the criminal IP addresses, statement methods and whatnot, but then immediately disrupting the DNS resolution such that they cannot use it as a renderable communication point. And the other forms or other operational approaches in order to start helping mitigation is to sinkhole those domain names.

And sinkhole really means is to register those domain names and use a control name server so that we can enable telemetry and do analysis of the traffic and try to pinpoint where is that communication coming from or going to, and also help identify infected systems so that you can go into the next two stages, which

I will not go into much detail because that's where the ICANN ecosystem, it's really not that effective.

We're talking about mitigating identifying infected systems, deploying patches to software in order to fix those issues in those systems, networks, endpoints, and what have you. So, the key takeaway that I want you to get from this illustration is that when we're talking about an effective response playbook, it's more than registry operators or registrars or ICANN itself that play a part here. We are part of a wider ecosystem. And so, the goal for any policy work that we want to do, and I'm going to talk about that in a minute, is to help strengthen the whole ecosystem such that the whole group can work together. So, let's go to the next slide.

And in order to illustrate these concepts about what the DGA, the asymmetry issue with the thousands of domain names that can potentially be created by these DGAs and how the actors that need to work together in order to mitigate these threats, we're going to be looking at a case study, which is a very well-known DGA that happened many years ago, but it's a well-known, well-studied, which is the Conficker. We'll touch on the thread profile, what was the response to that Conficker threat, and what were some of the lessons learned, and some of those are all inclusive use today. So, I'll get to that in a second.

So, let's go to the next slide, please. So, what was Conficker? Conficker was a DGA botnet malware that dated back in 2008. I'm not going to get into the technical details here and there. But what

you need to know here, Conficker went through different iterations, and they were known as the Valiant A, B and C. And what you see here in big, bold numbers is that Valiant A, as when the Conficker malware was first released, it created the DGA attached to that malware had created 250 domain names a day across five top-level domain names.

Once that was detected, it evolved and created the Variant B, and that stayed in the same range of 250 domain names a day, but then it increased the spread of TLDs to eight. By the time the malware morphed into the Variant C, we were talking about 50,000 domain names a day across more than 100 TLDs. And this is where it really gets into the it's more difficult operationally and financially to start blocking these domain names. And you can think now, when we now have thousands of TLDs, that issue can escalate very, very quickly.

Next slide, please. And so, in order to address the Conficker threat, a group of volunteers, which happened to be known as the Conficker Working Group, just worked together. ICANN was very instrumental in helping facilitate the coordination between cybersecurity researchers, law enforcement with the registry operators, both in the gTLD world and the ccTLD world. And the key takeaway here, again, is the number of actors that need to play in a collaborative way in order to effectively address, disrupt, mitigate, and remediate the threat.

Not just blocking a DGA from not registering those domain names, that's the end of the story. Really, it's just one of the steps that need to happen here. But we can only work the scope of our own capabilities which is the DNS, TLDs, and second layer domain names.

Let's go to the next slide, please. And as far as lessons learned, and again, we're talking about a period of between 2008 and a few years down the road, the critical lessons that were determined back in the day is that registry-level block lists do not scale. As you can see, by the time Conficker was in the Variant C, we're talking about 50,000 domain names across 100 TLDs. That really adds up and it starts to be a whack-a-mole problem. And so, really not sustainable in the long run. And that's how operational approaches in terms that the registry operators can deploy in order to block the communications to those command control domain names has also evolved.

There's a range of options that we can use. Another lesson there or tool that was established or created back in the day in order to help registry operators do what they need to do in order to disrupt these DGA domain names was the creation of the expedited registry security request. Now, it is known as a security response waiver.

But that was actually one of the good items that was established by the lessons learned by the Conficker Working Group. Next slide, please. And I believe John Crain is with us in the room. John, I don't want to put you on the spot, but if you want to share any insights

or perspective about how ICANN plays a role in the Conficker Working Group, please do chime in.

JOHN CRAIN

Thank you. We're going back in history. 16, 17 years now, so it was quite a while ago. I think, not ICANN, but the community that got involved in that was trying to solve a problem we hadn't seen before. I probably shouldn't say this, but we were kind of making it up as we went along. I think there was a lot of good lessons learned in Conficker.

I think there was even better lessons learned during Avalanche and some of the following takedowns. I think you know we learned a lot. We didn't necessarily get all those lessons out there into the community as well as maybe we could have done. Not all of the things that were suggested for improvement were really taken up by the community. I saw your mention of DNS cert, that died on the vine.

I think it was before its time. This idea that we need to communicate better, we need to automate more things, I think the concept was right, the timing was wrong, and I'm seeing a lot of discussion coming out of the community about how do we solve those same problems. It was fun, it was exciting times, we were trying to change the way we did things, and I think we had a lot of positive outcomes and I think we also realized that you know we'd

never get a perfect solution to this, but you should never let the perfect be the enemy of the good.

And I think one of the things I also realized over the years is a lot of the thoughts on the industry side, it was more about blocking communication because we're DNS guys. But on the law enforcement and the security practitioner side, it was probably more about victim mitigation and being able to get information so they could actually reach out to victims, which is not ICANN's role. But the industry, I think, stepped up and played its part. I think overall it was a good thing.

DENNIS TAN

Thank you, John. Yes, indeed, Conficker and then Avalanche taught us a few items here and there that we potentially want to talk about in this time now. And this is exactly what this is intended as far as the issue report and being a priority issue. We, the CPH and Registry Stakeholder Group as well, we see an opportunity to have targeted improvements to the existing framework.

And again, we can revisit the history, history can be a good teacher, and then from there, we can, you know, move along. So, in that spirit, what is it that we have today? And this is just a fraction of what we have, the elements today. Right here is we now focus on, we talk about the ecosystem, the range of players that take into account that have to work together in order to effectively remediate a DGA threat.

But now, we're zooming into what we can do here at ICANN as registry operators. And using one of the examples that we've produced, in 2021, joined with the Public Safety Working Group from the GAC, and I believe Gabe is with us in the room, so we work together in order to create this resource for law enforcement and registry operators in order to so that they understand what are the options we register operators have in terms to disrupt communication of these domain names. But also serves a guide for law enforcement to help in communications and how to deal with these domain names.

Again, there are a range of options in terms of how we can deal with trying to disrupt communication. And one of those, I repeat, reservation of a domain name, preemptive registration, disrupt the resolution, sinkholing, and depends on what we need to do. And John mentioned about remediation and trying to enable telemetry so that we can analyze and identify infected systems, get in touch with those systems and try to remove the malware and whatnot and do the whole thing complete.

So, this is where we are today. And what we see is that we can improve upon the system framework because this might work for maybe a small scale or some medium scale threats that pertains to DGAs and malware and botnets, but when the next Conficker happens, when Avalanche happens, can we have something in the ready to act in a coordinated and efficient way when, again, that threat level happens? So, let's go to the next slide, please.

And this is where we believe we publish a white paper along with the comments to the issue report. And we propose assumptions for the next policy work, and these assumptions are based on you know targeted improvements based on what we have today that centers into create a facilitation role. ICANN might be one who can play that role, facilitate the communications with those in the cyber security law enforcement with the TLD operators.

And when I say TLD operators, I'm talking about the gTLDs and ccTLDs, because DGAs are agnostic of TLDs, they can affect, they can work, they can use or misuse both TLDs and ccTLDs with disregarding who operates what. We also see potentially an opportunity to see how, now that we have the security response waiver, and that works one-on-one, every time a registry operator needs to utilize the security response waiver, can do so, and there are avenues to do that.

But in the case where the scale of the threat is such of a size, can we find different ways to act on these security response waivers in a way that is more efficient. Instead of each registrar doing that, what if there is a different mechanism in order to provide those waivers in a more efficient way across the board? Last but not least, the disposition of DGA domain names.

One thing is to disrupt the communication, but another thing is also to enable telemetry and provide information for remediation work. So, can we also talk about ways in that we tell the TLD operators can help enable that kind of function? For example, the

sinkholing function, right? There are certain registrars that specialize in that kind of function, but can we talk about sinkholing as a function rather than as a business model, if you will?

Also, potentially, we TLD operators can improve signals when we act on domain names that are sponsored by a registrar so that registrars understand what's happening with the domain name. So, we are talking about the technical level, what kind of signals we can provide our registrars, colleagues, or friends so that they understand when a domain name is acted upon, based on some security threat, they understand what's happening so that they can provide better information to their customers, but also enable communications back to the registry operator as well.

And last but not least, registrar remedies. And this pertains to when and this might happen with core orders when domain names are acting upon and they need to be blocked or suspended in a way but the registrars still hold the domain names, and so what happens in that in those situations, and that has happened in the past as well. So, again, these are our contributions or ideas as to how we think we can do, again, targeted improvements, marginal improvements upon the existing framework.

So, we look forward to continue this conversation throughout the weeks, months that we have in order to refine the scope of the work that we have in line. Could be a PDP to solve certain things, but we can also find other avenues to address other issues as well, right? Could be best practices, all technical work in the case, for example,

signals from registry to registrar. Let's go to the next slide. Okay, so now we go into the panel conversation.

I think I'm going to invite John Crain and Gabe Andrews to give us your perspective in terms of challenges in the current security threat landscape as it pertains to botnet command and control DGAs. Please also weigh in on measuring DGA mitigation. How can ICANN policy work might enable a no measuring DGA mitigation, and also how can we enable potentially remediation approaches? So, let's go around. Gabe, can I start with you, and then John?

GABRIEL ANDREWS Sure thing, Dennis. Hi, guys. This is Gabriel Andrews, co-chair of the Public Safety Working Group. And before I give some thoughts, how are we doing on time? I don't want to talk a lot if we're tight.

DENNIS TAN We're good.

GABIREL ANDREWS Okay. Well, then yeah, I think you and John have hit a lot of the very same points I would have brought up. So, first of all thank you for walking through this. I'll just share a little bit of the flavor that's purely from the law enforcement perspective then, And I guess just to reiterate, this is not a, oh, I see a note, Gabe needs to speak into the mic, not just adjacent to it. Okay.

This is sometimes hard work, but it's hard work because that's the intent of the bad guys, right? They're using DGAs to make life hard on us, and that means this isn't going to be trivial to solve. Back in the old days, malware used to have a hard-coded IP address so that when the bad guy has a half a million compromised devices out there in the world, he could just tell them all, hey, find instructions at this IP.

We, cops, slow as we are, we caught on to that, and we can go out and seize that IP and stop it. And so, the next bad guy comes out and he says, ah, I'm smarter, I'm going to use a domain this time, and that domain can point anywhere, and I'll change where it points rapidly. And so, we go out, we seize the domain. And that's why they turned to these DGAs, because now they can just put a piece of code on those victim machines that cycles through rapidly and every five minutes checks different domain that might not even exist yet but will in the future that he can register and give instructions.

And what makes this really challenging is that you have to stop that from ever happening ever again throughout all infinite time going forward otherwise all it takes is one occasion for the bad guy to get that communication out and then he can give all the instructions to all those victims again and then you start all over. So, the example that I had, I was just speaking to the case agent that worked with John and others on the Avalanche botnet takedown, this is like Conficker, it was another DGA based botnet.

This was mostly the time frame between 2009 up until the remediation action was taken globally in 2016. We didn't start that case in the FBI in the US, it was started I think in Germany, but it highlights for us how important the collaborative effort is. So, we have law enforcement all around the world that all recognize like hey this is a really big deal that affects all of us.

The victims of these machines were getting hit with ransomware, bank fraud, and so forth, the machines were being used for phishing, all of the above. And we felt that we had to take global action on it, and it simply could not be done without a lot of collaboration with the ICANN space, with the registry space. And registries more than registrars because, again, you're talking about domains that haven't been registered yet.

The bad guy can go to any registrar in the world to register a domain. Thankfully, though, because we have really smart people in the private sector that can look at that malware and reverse engineer the DGA, we can figure out, hey, where is that domain name going to be? Or what domain is it going to be a year from now on Tuesday? And so, as long as you can stop that from actually being available to the bad guy on that day, that one narrow band of time, then you can stop those new instructions from going out to all the machines in the world that the bad guys are using to commit harm.

So, the challenge then is, how do you do that for all of the registry TLDs that might be listed within that DGA? And it's not your fault if

you happen to be one of the outputs, it's just whatever the bad guy chose. But how do you do it in a way that you can be sure that you aren't going to fail once. Because if you fail once, that's all it takes for the bad guys to reestablish control. So, I commend some of the effort that went into that framework for addressing botnets or DGA in use with botnets.

And I especially commend the effort to try to align the incentives because some of the pain points we found are that some registries might have wanted us as law enforcement to get court orders for all the action. That works for a little bit. With Avalanche, I think we did for a while. We went to our courts in the U.S. and we got seizure orders for hundreds of thousands of domains for like just one year, for the next year. These are all the domains that this is going to put out. We go to the courts, we get the permission, we give that to the registries, and then we have to do it again next year and then have to do it again next year for infinite time.

And this is one botnet in one country, and then if there's another botnet doing this, do you have to do it again? And at some point, it's just not scalable and it will overload the courts and so forth. And so, moving to a system where ICANN and registries recognize that this this is something that is important to do and will accept conversations about voluntary action to act on these DGAs is essential in order for this to be viable going forward.

Further it's essential that they don't have misaligned incentives like avoiding the fees that ICANN would otherwise put in place for

registries to reserve those domains. We don't want a registry to have to pay money to do the right thing. So, I think all this collaborative work to be flexible and to address it together as a community was essential and will continue to be essential.

And I think that's really what I would like to stress with this. And I very much appreciate the thoughts going forward about how to make it more systemic and to really enable this to be protective for as long as we have to face the threat, which I don't think is going away anytime soon. John?

JOHN CRAIN Well, that's depressing. I was hoping the bad guys would just stop doing bad things. So, reflecting a little bit on that, I think one of the things we really struggled with and still do a little bit is that law enforcement do not understand in general how this magic at the back of the internet works, they have no clue. Why would they? Even when you've reversed engineered the DGA, that's not the end of it. You now actually go and make sure that you're not causing more harm than good.

And the bad guys know this, so they put in live names, they put in names of things that are critical, and they try everything they can to obfuscate and make it harder for you to understand when it's safe to block a name or register a name, and when you actually need to take a name away from somebody, because that does

happen, and that's a lot of work. And there's also not generally an understanding of the intricacies of the industry.

So, often, and if you've come to ICANN meetings, you've seen people from outside the industry assume that everybody in the industry is way smarter than them and understands how all the technology works. And because they have an idea, it must therefore be easy to implement. So, while in Conficker and Avalanche, we had time-based DGAs that in theory you could register a name for 48 hours and block it, most of the people in the industry didn't have systems in place for that.

So, they do hand blocking of names and they just take the entire list and block them for the entirety of, well, forever basically. So, we have registries out there that have block lists that are hundreds of thousands of names long, and arguably those names don't need to be blocked anymore. So, this was not sophisticated.

Now, I don't know if we need to be more sophisticated or not because it comes with cost, but it is not as sophisticated as sometimes people think it is. They think there's a lot of sophisticated engineering here and that just isn't. And that comes a level of risk on the business side, if you like, on the registry and registrar side, because you've not done all of the evaluation. So, you have to be able to have this trust element. There are some fantastic organizations out there like Fraunhofer Shadowserver, Kumri [ph], who put a lot of effort into looking at these DGAs and doing the analysis.

But when the list comes over to a registry, how do they know that this has the right evidence, the right level of due diligence? We've never really talked about that. In the DNS abuse policy, we talk about providing evidence. If we're going to do something in the DGA realm, I would expect that there's going to be some discussion around that. What is the expected level of data that you're going to give to the registries and the registrars?

And I'll be honest, the trust isn't always there either. So if you're going to give a list of, I don't know, just say 10,000 names that are in the future, give you the ability to take control of a large piece of malware, that trust has to be there. You have to know that you're giving it to the right people. And none of that is really in place as an industry.

We put forward the DNS cert idea, like I said, way before its time, but we don't have those mechanisms. And I will point out that that's not a unique situation to our industry. That is across most industries that information sharing just turns out to be hard. So, if we want to make this easy, we've got a lot of hard work ahead of us.

DENNIS TAN Thank you, John. Thank you, Gabe, for those insights and perspective. Very useful. And now I think we have a couple of minutes for questions from the audience. Reg, you can help me moderate here if you see.

REG LEVY Absolutely.

DENNSIT AN Okay, I have a queue.

REG LEVY Go ahead.

CHRIS LEWIS-EVANS Yeah, thanks, Reg. Thanks, Dennis. John and Gabe, great points. Thank you very much. I worked a little bit with Gabe on the 2021 proposal and framework, and I really support it. It's all about the collaboration. I think one thing to note, there was a lot of talk about C2, so command and control, and botnets. The criminals have moved on a little bit, and they use DGAs for phishing and all sorts.

Rowena might even present, I don't know what she's presenting, but she's talked about the winter fuel problem for the UK, that was a DGA. That could have been blocked, some aspect of that. And how we make it easy, how we process that is really key. So, yeah, really looking forward to this work. I think it will have a big impact and we're able to do it both on the mitigation side and sort of helping people out. Thank you.

REG LEVY Thanks, Chris. Michele?

MICHELE NEYLON Good afternoon; Michele. This is probably a slightly dumb question in some respects, but would somebody be able to clarify for me whether the DGA domains are still mumbo-jumbo gibberish or are they something that could be semantically useful?

JOHN CRAIN The answer is, yes, they are both. And DGAs can be 100,000 names, but they can also be 10. So, the criminals are smart. The guys that law enforcement are tackling here and public safety folks, they're not stupid, so they deliberately use things with semantic meaning to make it harder to block it because obviously, that may have a real legitimate use.

How do you tell that a name that has been registered, what looks like a legitimate website on it is actually belonging to the guy who owns the DGA or not? So, these guys are very smart. So, yeah, it's both. There's a lot of gibberish out there, but there's also very clearly legitimate-looking names. And I'm not talking about the kind of stuff that's used in phish where it just kind of looks a legitimate. These are actually very legitimate looking names.

MICHELE NEYLON Yeah, thanks. I'm just following through on that. So, I think the issue then I think as already mentioned will be very much one around trust because obviously if you're telling either registries,

hey, all this stuff that you could have sold and made oodles of cash with, no, you can't. That's obviously a problem. But I doubt most of them were particularly interested in enabling scumbags to proliferate across the internet either, so trying to get that balance right.

I suppose the big thing when it comes to any of this is that whether the way to get there is through one means or another, there shouldn't be an issue around the finances coming from ICANN, in other words your employer, John. You guys need to be able to say clearly and categorically, we will facilitate you in doing this and we will not make you go through oodles of paperwork that keep legal departments in jobs.

JOHN CRAIN Well, I wouldn't want to make my lawyer colleagues unemployed. I think that's a little bit too much. But we did the security waiver trying to keep it as lightweight as we can. And to keep it as open as we can, we wrote it because of the DGAs, but our principle is that we do not want to get in the way of the industry doing the right thing. There are some safeguards in there so that we can tell that they're actually doing something with these names related to it.

And there's no policy really around this. This is just contractual. So, if we need to change that in a way that makes it just much easier for us to get out of everybody's way, which we would love, by the way, then maybe that's an area to think about in the policy realm.

I don't think we've ever turned down a waiver. We don't want the money. This is dirty money used for dirty things. We don't want it.

GABRIEL ANDREWS

Can I chime in briefly on your point about there being, and John's point, about there being domains that are intentionally real appearing? There's also, to his point, sometimes intentional collisions with pre-existing domains, and I think it's important to recognize that the mitigative action you take can be entirely different depending on whether the domain has never been registered, and it's just like this perspective registration in the future versus a domain that's been long established and that exists in the world.

It doesn't mean that you completely ignore that existing domain, but maybe you just do a little bit of extra of checking to make sure, okay, it's not the bad guy controlling it, and it's only that one, usually, that one point in time that that it will be potentially dangerous. That's an entirely different tactic to take for that that potential collision with another potentially innocent person than the completely not yet registered domain that you just want to make sure doesn't get registered at the wrong date and time.

REG LEVY

[00:46:43 – inaudible].

UNKNOWN SPEAKER

So, as you've mentioned, Taher, that there is a lot of hard work to do before it actually becomes easier, as DGAs has become more sophisticated potentially using AI or block chaining naming systems, do you think DNS based mitigation strategies like sink holing will still remain effective or will the community need to rethink the architecture of abuse response entirely?

JOHN CRAIN

You want to have a shot at that, Gabe, or do you want me to?

GABRIEL ANDREWS

I'll take just a quick shot. So, blockchain domains, typically, we're a little bit lucky that blockchain domains require browser plugins and so forth, they're not quite the same risk profile, but there's a lot of points that you raised there. There are a few key concepts that I think hold true that if we know what those domains are, we can do sink holing which means we control it for just a brief amount of time to see all the potential victims that are reaching out.

Well, we can make them reach out to us, and that's all that sink holing is. Instead of they reach out to the bad guys, they reach out to a good guy, and that enables us to then go out and notify all those victims that, hey, you've got a compromised machine on your network, your town, your business, your hospital, wherever that machine may sit. That's always going to be something that we're interested in doing, and it's a useful action to take.

But I think to your broader point that as the technological landscape evolves, absolutely there's an ongoing cat and mouse game. That's never going to stop. We're always going to have to keep fighting against the bad guys. And they're always going to be at least as smart as we are, to John's point. I'm getting older. I'm not getting younger. It's work.

JOHN CRAIN I'm getting younger. And don't forget the tools, this has been going on for decades, the tools that enable the bad guys to do more things or things faster, are the same tools that we can also use. So, the question is how are we going to use those? And I don't mean we, ICANN, I mean with my sort of OPSEC, Operational Security hat on, how are we going to use those tools?

And there is a broad discussion going on in the operational security world at this moment about exactly that. Because it's not just for DGAs, it's for all online harms. The bad guys will use whatever tools they can get, and they're well organized and they don't need to necessarily have the high tech. They can buy a lot of these tools, but we can also use tools. So, it's not all everybody likes to do the doom and gloom about emerging technology and have ever since the steam engine. But we can use steam engines too.

REG LEVY There's no one else in the queue, Dennis. Back to you.

DENNIS TAN Thank you, Reg. Great question. Great conversation. Thank you very much. Now at this point in time, we'll go into part two of our agenda. So, Rowena Schoo, please.

REG LEVY There is a hand in the room. Go ahead.

NIKHIL MATHUR Hello everyone, my name is Nikhil Mathur. I am a NextGen Fellow. My question was related, how can PDP 1's associated domain name checks leverage DGA high entropy patterns as triggers to disrupt tunneling enabled botnet portfolios?

The reason I am raising this question is inference from the INFERMAL report that that the ICANN often discusses that repeatedly mentions the C2 botnet exfiltration. And the second thing I'm curious to know that while checking the malicious domains, theories like Shannon's theorem can be practically used in to verify the randomness of the domain names. Thank you.

DENNIS TAN Thank you, Nikhil, for the question. I can take it. I think at a high level, you mentioned the relationship between the current PDP on associated domain check and the work that we can do here on the DGA. So, at the fundamental level, when we talk about associated domain check, we're talking about registered domain names. And

when we're talking about DGAs, it's mostly unregistered domain names that we want to block in the future.

So, there is no relationship at the level because associated domain check is a targeted PDP looking at when a registrar is investigating a confirmed case of DNS abuse in a registered domain name, then they will pivot to look at other domain names, other registered domain names in the same account or associated to that domain name. So, there's the difference between ADC and what we can do for DGAs. So, hopefully that answered your question. Okay, so now -- sorry, Reg.

REG LEVY I was just going to turn it back over to Rowena.

DENNIS TAN Oh yes, let's go. Rowena, please go ahead.

ROWENA SCHOO Thank you so much, everyone. Alright, I'm not going to talk about winter fuel. I don't know if Chris is still here, but I will be talking about that at the Contracted Parties Summit. So, today we're going to talk about how mitigation is measured, and I'm going to walk you through some of our thinking about how we attribute that. I'm Rowena Schoo from the NetBeacon Institute. If we can go next slide. Thanks.

I'll tell you a little bit about what the institute is, give you some understanding of the project that we use for this measurement, talk a little bit about the different forms of phishing that we see, and then I'm going to do really a deep dive on how we measure that and how we might think about attributing it to specific entities.

We get a lot of suggestions and requests that we publish more information around mitigation, and we've been working towards that for some time, so I want to use this as an opportunity to share some of that thinking with the community. Next slide.

So, the NetBeacon Institute is part of the Public Interest Registry. We were created in pursuit of the non-profit mission that PIR has. They are a charity in the US, so they reinvest in issues that are of public purpose, and we are one of those. But what we do is much more outward focus. We don't work on the registry side. Internally, we create tools and services to try to help with DNS abuse. And that means that everything we talk about is non-commercial.

We're not charging for this. It's free. Everything we're offering is not associated with the charge. We're not selling anything. So, what we do is a selection of things. We do a lot of work around data and analysis and best practices. We do a lot of outreach and education and we work across the ICANN community but also beyond that as well. We're really trying to create a safer internet for everyone. Next slide.

So, you can kind of group what we do into three pillars. Generally, we try to understand DNS abuse, we try to disrupt it and we try to

prevent it. I'm really going to be focusing on the first pillar here today and talking a lot about our project NetBeacon Map, so the NetBeacon Measurement and Analytics Platform. But we have a bunch of other resources on our website if you're interested in understanding these issues further. Next slide.

The NetBeacon Map is a project that we run in collaboration with our research partner, which is KOR Labs. They're based out of the Grenoble University in France. And they do all of the methodology and the academic collection of the data and analysis that sits behind this. The benefit of working with KOR Labs is that they are independent, they have a credible academic methodology for doing this information, and so we can find that the data they give us is very, very reliable over time and for comparing across the ecosystem.

We then package this information up into a selection of different resources. So, we publish publicly a lot of information in our interactive charts and our monthly analysis, but we also have individual dashboards that we provide for registries and registrars that gives them detailed information about phishing and malware in their zone. And I should say, when we started this project, we really wanted to measure the whole category of DNS abuse as it's defined within the ICANN space.

But what we found working with KOR Labs is that they felt at the moment they could really do this reliably and accurately for phishing and malware, but it was more difficult for the other types

of harm. So, everything I talked about today is just focusing on the use of the DNS for phishing and malware. Next slide. So, if you really want to dive into our methodology, it's that document on the left-hand image that's available on our website. Hi, JC. No worries.

That's going to go into more detail about exactly how the project works. If you prefer a more of a summary approach, we've also got a report that I've pulled together which is where the image on the right comes from. That gives you a nice visual of how this is put together. I will say we start from a similar place to most measurement projects in the sense that we start from an imperfect place using threat feeds that are typically created for network protection, not DNS level mitigation.

So, KOR Labs have picked four that they find to be the most reliable with the least amount of false positives. From that point, they do a huge amount of analysis over the top of that. So, some of the key features that I'll call out, they exclude what they call special domain names. So, this is domains that are never really going to be appropriate for DNS level mitigation. So, subdomain providers, dynamic DNS, file hosting sites, things that are being sort of abused as a service rather than being appropriate for DNS level suspension.

They also have a process which is quite complicated. It's a proprietary process that they've developed that classifies the domain names into whether they are maliciously registered, so for the purpose of phishing and malware, or whether they are a case of compromise. So, usually that's going to be at the website level.

And this allows us to do a lot in the analysis, which is really helpful for understanding the issue.

So, I think we can go to the next slide. If you want more detail on the mitigation, sorry, on the measurement, definitely have a look at those documents. And let's talk a little bit about what makes a good mitigation measure.

Next slide. So, this is a framework that we use at the Institute to think about mitigating harm online. And we're really talking about like what does ideal mitigation look like? What attributes would you want if you were mitigating harm? And we think that looks something like this. We want the action taken to be effective, we want it to be quick, we want it to be relatively simple ideally, we want it to be precise and proportionate to the harm that we're seeing so you're not having a huge amount of collateral damage as a result. We want it to be cost effective and also necessary as a mechanism for disrupting this harm.

And I think when we think about mitigating at the DNS level, often we see that it's on the element of being precise and being proportionate that there are challenges and that the DNS is not the ideal mechanism for mitigating some types of harms.

So, if we go to the next slide. When we think about what actions could be taken, so if a registrar gets a report of abuse, there's really only a limited amount of tools that they have to be able to mitigate that harm. If the domain has been registered recently, they could delete it. That might be preferable if it's in the AGB period.

Otherwise, they've got the option of suspension, which means the domain is not going to resolve, but the content will still exist.

Alternatively, we've talked a little bit about sinkholing, that is also an option that sometimes gets used or they can refer it to the registrar or the web host or both. Sometimes they're the same, sometimes there's a reseller involved. And if you're looking at a registry, they would also have similar kinds of options but also be able to refer it to the registrar. Next slide.

So, just one important note, I mentioned the distinction between malicious registrations and cases of compromise. I just wanted to share some examples of what that would look like. The example in blue here is an example of a malicious registration where the whole domain name is being used for the purpose of the phishing. In this case, it's pretending to be a government website. The second example in green is what it would look like if there was a compromise in the website. So, this domain name on the in the green text, you've got an unrelated domain name but then after that you've got a page path which is going to what is trying to collect PayPal credentials.

So, something that is definitely trying to collect sensitive information, but it's not the whole domain name that is dedicated to that. Typically, when we see compromise, it's happening at a content management system that's had a vulnerability that has been exposed and has been accessed by someone who's acting maliciously rather than the registrant acting maliciously. And this

is important because suspending the whole domain name is going to have a lot of collateral damage and disrupt the email of the victim who is really the registrant. Okay, next slide.

So, with all of that in mind, I wanted to talk about how we think about measuring mitigation. So, next slide. So, for a long time, we've been publishing charts on our website that give how much of -- sorry, I should start again. When we measure phishing and malware being used in the DNS, we also measure whether mitigation is happening through this project. And so, we've been publishing charts that do this at an aggregate level. They're already there today you can have a look at them.

But what we've also added recently is this drop down on the right where you can isolate for malicious registrations or compromised. So, that's going to change the numbers slightly. You can also, now on the left, change from looking at the count of domain names, which is what you're seeing now, so that you're looking at a percentage. If we go to the next slide.

So, that would be the percentage view, which is maybe easier to track over time what is happening. These types of charts we also give to registrars and registries in their dashboards, and I think the really important thing to note about how we measure mitigation at the moment is that we don't attribute it to specific entities. We are really just looking to see if the harm has stopped, and that could have been a selection of different actions.

When we provide it to a registry or a registrar, what we're saying is that this is the mitigation rate for the domains in your zone. We're not saying that definitely it was them that took that action. It could have been a variety of actors. So, if we go to the next slide. So, we're really looking to see has the harm stopped, and there's various signals of this that we use to understand it. I've tried to set out a selection of those on the screen.

This is not comprehensive, but I hope it gives you a bit of the picture and a flavor for how different actors can take different action that would show is mitigated in our system. In the gTLD space, there's a differentiation between server hold and client hold with server hold being applied by the registry, client hold applied by the registrar, but in some cases if there's a reseller, they might also be able to apply client hold. In the CC space, this typically has different names and is treated in slightly different ways, but there's usually a way to suspend or disable the domain name.

Our mitigation metric at the moment will also include actions from other entities outside of registries and registrars. So, if the host acts and removes content or suspends a hosting account, that's going to show up as mitigated as well. It could also include the registrar not acting or changing things or other actions that we've been talking about today, like sinkholing, LEA seizures, splash pages, that type of stuff. And this is quite important because this complex ecosystem is happening all the time.

And when there is a report of abuse, it's quite common that there's referrals between these different entities, and maybe one of them is acting, maybe multiples are acting at the same time or shortly after each other.

Next slide. So, we also measure mitigation at different times throughout the month. So, when a domain comes through one of our source feeds, it ends up being tracked for 30 days. We take these measurements at smaller intervals to start with, and then they extend out into hours, and then finally every 12 hours for 30 days. And so, each month exists independently. We don't kind of have a cumulative carryover between the months. It's just within that month we're checking for mitigation for that 30-day period.

I think I'll go next slide. Okay, so some of the challenges that we see in this, it's pretty hard to actually attribute to specific actors. We provide this information for registries and registrars. We've got reports where we publish abuse rates, and it's quite a common request that people want to see mitigation rates in there as well, which we would like to publish.

But I think that there is probably a logical conclusion that if we were to publish a mitigation rate next to the name of a registry or registrar, people would think that we were saying that specific entity took that action. And that's a pretty tricky problem to solve in terms of how you attribute that busy ecosystem to specific entities. It gets even more complicated when you have various

evasion techniques that can make measurement hard, that could include things like geo-blocking.

I mentioned at the start that the source feeds are not perfect, so sometimes we do see false positives which cause complications, or we see different types of harms come through that don't actually fit within the definition of phishing and malware. For example, fake web shops. We've also had instances of incomplete URLs coming through the system, which can make it hard to measure mitigation accurately because of the way that we're looking for different types of signals including content changes. Okay. Next slide.

So, I guess to wrap up I wanted to give you a sense of how we're thinking about this in terms of attribution. I think if you take those signals that I put on the screen, we can kind of bucket them into attribution to particular entities remembering that in some cases multiple entities will be acting at different times across the month. In the cases where we can see client hold, we can kind of attribute that to a registrar, although noting it could potentially be a reseller.

Sometimes we'll see client hold and server hold that can look like a registrar and registry, server hold at the registry level, but again sometimes those detailed attribution elements are not available either due to the way that it's measured or due to it being a ccTLD, in which case we might be able to see that something has happened at the DNS level, but we can't specifically say if it was the registry or registrar, it could have been either.

And then we have this bucket which is various, so I know that's not very satisfying when we're talking about attribution, but this is going to capture some of those other actions that could happen by a variety of other actors. So, that includes if the host acts and changes the content, the sinkholing, LEA seizures, those types of signals that are more difficult to pin down to specific entities because they could happen at multiple stages.

And then finally, I mentioned some of the complications in terms of measurement. We're thinking some of the responses we get are actually a little bit inconclusive, and that bucket would be intended to cover some of those complications some of the evasion techniques that we see. All right, next slide.

So, I think basically summing up, I wanted to really share some of the complexity that we're trying to measure here around mitigation and the different ways that we're thinking about that, some of the challenges we're seeing, and our initial thoughts on how we might attribute that between different entities with the intention to bring more data out to you soon. I think I will stop there and take any questions.

RG LEVY Don't be shy. Michele, go ahead.

MICHELE NEYLON Yeah, thanks; Michele. I mean, it could just be that your slides were so good, Rowena. Did you answer all of our questions

preemptively? No, I'm joking. Well, not really, actually. There were good slides. I think this is good for us explaining a lot of the landscape to people, which is really, really important. Because I've heard, again, in many places, I mean, John intimated it earlier, people come along and go, hey, I've got this really, really cool idea for fixing random thing.

And then when you try to explain to them that, no, it's not that simple because that's not how that stuff really works and trying to get them to understand there's layers to it, we're not being awkward, we just can't do what we can't do. It is helpful. Thank you.

ROWENA SCHOO

Thanks, Michele. And I think even as I've worked on this quite a lot still coming here, I still had conversations where I'm like, oh, is that how it works for this particular circumstance? It's only recently that it clicked for me that a reseller could also put client hold on or that a reseller is probably going to be the host.

So, actually that's probably going to show up in that various category in our data which is again somewhat unsatisfying because we don't know who did what exactly. But we're trying to work through this and talk to people as much as possible while we're here to really understand what these signals are telling us just.

MICHELE NEYLON: I'm going to throw a real kind of curveball at you. In terms of when you're checking to see whether the content, site, whatever you're calling it, is still up, are you doing anything with device-specific checks? So, what we've seen, for example, is where we are seeing things where, first off, they're blocking my entire AS number or my primary AS, so we can't see a damn thing.

So, we have no way of verifying anything. And the other one is where the content will only appear if you use a specific device. If you try anything else, it won't work. So, that explanation was the rest of the room, not for Rowena who understands this, I think.

ROWENA SCHOO: Yeah, so that is a real issue and a challenge when we're doing this. And when I say we, it's really KOR Labs doing all the complicated stuff. I think what you're talking about, I think, would fall into some of the inconclusive categories where it's like geo-blocking or it's only available on mobile in a specific country. And that is a real challenge to try to figure out exactly what's going on. Yeah, it's complicated.

REG LEVY: Thanks, Rowena, and thank you, everyone, for your diligent attention. I would like to invite everyone after the coffee break to come join us in room 1:05 for more DNS abuse. We're going to be doing CSI ICANN.

NIKKI SCHULER We can end recording. Thank you.

[END OF TRANSCRIPTION]