ICANN81 | AGM – DNS Abuse Updates
Wednesday, November 13, 2024 – 13:15 to 14:30 TRT

MERT SAKA: Hello, everyone. Welcome to the domain name system DNS abuse update session. My name is Mert Saka and I'll be monitoring this meeting for the chat room. In this role, I'm the voice of the remote participants, ensuring that they are heard and equally and those who are in the room participants.

Please note that the audio may be available in other UN languages and you can see the message in the text chat room. All chat sessions are being archived and follows the ICANN expected standards of behavior. I think with that, I think we can start, right?

MUKESH CHULANI: Thank you. Thank you, Mert. Welcome, everyone. My name is Mukesh Chulani. I'm from ICANN global domains and strategy division. We're really happy to have a packed house today to go through our recent updates, focusing on the updates of our office of the CTO. If you can advance the slide, Mert.

So we're going to cover updates on domain Metrica, and that'll be covered by Samaneh and Siôn Lloyd here. And then we'll have updates on INFERMAL. Maciej is here on my right to cover that topic from Kor Labs. And then just a few closing remarks, opening it to questions and answers from you all. Next slide, please, Mert.

We're going to go through this. I guess just to set the context for what we're about to see. We moved and have enacted the amendments for DNS abuse. So those have been applicable now for some seven odd months. You've seen the reports from contractual compliance and the initial findings. You're starting to see case level updates on breaches and actions being taken. That gives us a very narrow slice of visibility into cases being reported into contractual compliance. But there is a need for the wider angle view. And we have industry here as well who is doing that work and who are doing that work quite well. What we have here is some work being done by our Octo team to push forward this thinking, to have wider angle lens view of evolution of domain abuse. So I can turn over to Samaneh and Siôn to cover this topic. Thank you.

SAMANEH TAJALIZADEHKHOOB: Thanks, Mikesh. Hi, everybody. It's exciting and a bit stressful to see a packed house, but we will do our best. I'm Samaneh Tajalizadehkhoob. And today we will do this presentation with Dr. Siôn Lloyd. Both of us are from the Octo team, subgroup SSR research team. So I think it was about six years ago that I did my first presentation. I took over the DAR project and stood here also very stressed and talked about DAR. It's been six years now. We've received a lot of positive and negative feedback. We sat with several of you into sessions, improvements, comments, back and forth. Apart from what we did, others also developed tools out there, very good tools that are used for the same purpose, even better and better from our community. Tools like NetBeacon, CleanDNS, Adapt.Live from this community, but also

other tools from the academic community outside. What happened is that it was time to move on and to have a better platform that meets the needs of our community. This is what we are here to talk about, the ICANN domain Metrica measurement platform. Next slide, please.

Before going there, I will do a quick recap of who we are. So basically, the Octo team consists of four teams, research, technical engagement, SSR research and operations. We sit with the SSR research team. I lead the team and I work with Siôn and two other team members. Next slide, please. Next slide. Yes. Carlos, Sam and Siôn are members of the team. Next slide, please.

So this is the zest of what we do. We see ourselves as a think tank that work on research and development of Internet identifiers. We use data and scientific methods to answer questions that are interests of our community. Our goal is either to improve the existing metrics that are there to measure security or insecurity of identifiers or to develop new ones from scratch. Next slide, please.

This basically sums up the framework behind all of our work. We are interested to see what causes security or insecurity and how these controls that are in place to make sure that security is happening end up becoming attacks and attacks becoming incidents. So what we do in this group is to basically measure each point in this plot that you are doing.

In Metrica, we focus on the first module is about DNS abuse. I'll give you an example. There are two metrics that can be explained here. One is where is abuse concentrated. For example, you can look at the

reported counts, the number of abuse counts. And then what is a report? A report is a domain that shows up on a reputation block list. What's the difference between an attack? An attack is an attempt of using that reported domain to cause a damage. What is an incident? Is that attack actually materialized? Other measures are reactive measures that are takedowns, things that registrars and registries and hosting providers do to basically mitigate what they see on these RBLs. Proactive measures are controls, policies, configurations, things that they do before they see these domains on RBLs. Now, this whole session, Metrica and Infermo, will be two parts of this what I'm showing on this slide. The first part is what Metrica is focused. The second part is what Infermo measured. Next slide, please.

Now, going into the depth of the Metrica platform. Next slide, please. So it's slightly different from the systems that are out there in a sense that this is meant to be a measurement platform that has a dynamic dashboard and shows statistics. It's meant to be a home for data collection, combination of data, and giving that data to the community. Next slide, please. This is the idea behind Metrica for us. Actually, we gather input data. This could be RBLs, zone files, DNS records, email data, but also other non-abuse-related sources. We use the science that we developed inside SSR on various factors that would make metrics better, and then we create metrics. Next slide, please.

ICANN Domain Metrica is designed to have certain attributes. This is very important for us because this makes one of the main differences between this and R. These attributes design principles are based on the feedback that we've got for years from our community. So Metrica is

meant to be modular. That means that each module represents a metric, and modules will be added over time. So maybe each couple of years, one module, or each unit of time, one module, but it's meant to be home to our research results. So every research that we do, we make a prototype and add it to Metrica. It's meant to be flexible in a way that builds with requirements that will allow us to add these modules, and it's transparent in a way that we have control over it. We give control to the community on it. Methodologies and sources are public, and everything is meant to be shared with the community where our licenses allow.

Who are the teams that are working on this project? So we have the ICANN SSR research team providing the vision and the science, and we are working closely, we are engineering an IT team within the ICANN org to do the development. They provide the development resources. Why are we doing this? I think I already talked into that. With that, I hand into Dr. Siôn Lloyde, who will be the lead of the first module.

SIÔN LLOYD: Thank you. So this slide is the same one from above, but highlighted in green here is the areas of that which are covered within Metrica. So Module 1, as we have it so far, concentrates on looking at DNS abuse, looking at the reports that we receive within RBLs, and some decoration we can do around that, metadata, other things we collect, for example, geolocation of the abuse, and looking at various ways we can display that data so that it hopefully makes sense and is useful to people. Next slide, please.

So for Module 1, we have certain things in place, such as the statistics within just the gTLDs at the moment. We're not looking at ccTLD registrations. And you can search for things. So you can search for a TLD for a registrar or for a domain. And based on that search, you can see certain information, metadata around it, and bits of information that we've collected. We have shareable domain-level data, and that's just to the registries and the registrars who are responsible for those domains. So we can share lists within certain restrictions of the licensing, but we can provide that data onward so that people can see which domains are actually contributing to the statistics that we're displaying.

We have a user interface, so you can log in. You can have some interaction with the data. You can apply filters for things you want to see and things you don't want to see. You can add things, so you can compare multiple things together. We have the abuse map, which is using MaxMind data to geolocate where things are hosted. And we also have some popularity ranking for individual domains, so you can see how that popularity has changed over time. Alongside the user interface, there are other ways to access the data. So for registry and registrar users, there's the MoSAPI , and from there, you can get access to the data as it pertains to your particular entities. And we also have what we call the community API, which has a lot of the data that's behind the user interface, the more public side of the data as well. Next slide, please.

So where are we right now? In September, we launched a pilot release to a small group of users. The intention was to get feedback, so we

deliberately chose users who we thought would be very vocal and critical in their assessment of what we were providing, and that gave us lots of feedback that we could work with. And then in October, we released to the wider registry and registrar community. With MoSAPI following a day later. Now, the intention is to have a general launch in December, so that is to the wider ICANN community, anyone who's not part of the registry, registrar community.

We've put subject to change there, because where we are right now, we're responding to the feedback that we've gathered from the pilot phase and from the early use. Some of that feedback was about spam. Basically, the definition of spam, as we had used, was too broad. It was the same definition that we use in DA, and it doesn't match that in the contracts. So to avoid confusion in the short term, we're going to address that by removing spam from the system. The idea is that we have processes, we have research in place. We've published an OctoDoc 38, which looks at how we can more carefully categorize spam reports, so that we know we're looking at the more serious end of the spectrum. And when we've got that productionized, we can reintroduce spam to the system, but it'll be a much more constrained definition.

We initially had a few performance issues as well, so we're addressing bottlenecks, looking where we need to make things quicker, more responsive, because no one wants a UI that isn't reacting to clicks and so on. And we also know that a lot of the documentation wasn't accurate enough or comprehensive enough, I should say. So we're addressing that. It's hard to write FAQs before you have users on the

system, because you don't know what questions you're going to get asked. Now we've got those questions, we can address them more reasonably. Next slide, thank you.

So here we see, this is a screenshot of what you see when you initially log into the system. This comprises of kind of a general view of all of our data as we see it. So statistics for the gTLDs that we have access to, and the registrars. So we have numbers of aggregate volumes and how many things we've seen. Next slide, please. And we also have charts showing the history over time. So 12-month trends, or you can look at different time frames in other places. Next slide, please. And this is just to show that the UI has been designed that it also functions on small screen devices. So mobile and tablet, you get a slightly different experience, but it is a quality experience still. And you can still see all of the same data, make all of the same interactions that you can see on a desktop. Next slide, please.

So I said searchable you can search for things. As I said, TLDs, you can search for registrars, and you can search for individual domains. And then you get information specific to the item that you've just searched for. Part of that information is metadata. So, for example, if you search for a domain, you'll see the created date. You'll see the name service it's on, the IP addresses that it's resolved to, whether it's DNSSEC signed. For TLDs, you get a different set of data. You get the size of the TLD, the number of change in size from the previous measurement that we made, various different bits. Next slide, please.

So if you log in as a registry or a registrar user, you have access to the list of domains that have been reported for your entity, be that a particular TLD, a particular registrar account. And this is limited by some of the license agreements we have in place. But as far as we are able to share this data, we will share it. So you see the domain. You see which RBL it came from. You see what it was reported for. And in some cases, you might see how many entries it had on that list. So where there's multiple URLs on a particular domain, you'll see a URL count just to tell you how many times it's appeared. And this list can be filtered. So if you're only interested in malware, you can just look at that particular abuse type. If you're only interested in a particular RBL, you can filter on that as well. Next slide, please.

So one of the things we wanted to have is more interactions with the data. So where you see a chart, you can click into it. And then you can change the time frame that's being used--

MUKESH CHULANI: Okay, to make up for lost time, Siôn will speak at one and a half speed.

SIÔN LLOYD: Thank you. Apologies for the delay. So as I was saying, yes, here's one way you can compare different entities. This is a trend over time of what we have here is four different abuse types, spam, malware, phishing, and botnet C2. As I said, we will be removing spam in the short term because-- thank you. Next slide, please.

And here's actually that same data, just visualized differently. So now it's a proportion of the total within each entity rather than a trend across time. Next slide, please. And the lists of domains, so if you search for a TLD or a registrar and you see more than one domain, we can resolve those domains and geolocate the IP addresses so you can see where this stuff is being hosted. Again, this is broken down by the currently four will be three different abuse types separately. Next slide, please.

The other thing we have in the system is not abuse related. We look at the popularity. So here we see a search on a specific domain, example.com. You can see the metadata we talked about at the top. If it had any reports, it would appear in that middle slide. And then the third tile there is the Tranco popularity. And that graph we can click into. Next slide, please. So we can open an interactive version of that chart. We can change the time scale. And we can also add other entities into the comparison. Next slide, please.

So here we see the Tranco ranking over time. Forexample.com and icann.org as a comparison. And within this chart, you can also in the top right, you can download the data that's underlying it in either CSV or JSON format. So if you want to take this into your own systems and do something with it, but you don't want to interact with a full API, you can just do that quite simply. Set up the charts you want, get the data as you want to see it, and then download it and use it yourself. Next slide, please.

So what does this system have that adds to what we didn't have previously? We've got this interactive dashboard that you can then filter on. You can narrow down searches to look at things that are specifically interesting to you. Or if you see something that you think might be interesting, you can focus in on that. You can get the underlying data and then do your own analysis on it. You're not limited to what we offer within the user interface. We provide all the metadata. And the other thing that we're doing for this data is we're removing reports whether the domain in question is not within the zone file. So things that have been suspended and removed from the zone file won't be showing up in any of this data. They won't be contributing to the statistics. They won't be appearing in feeds. We've talked about the underlying domain names where we're allowed and the Tranco popularity. Next slide, please.

So I think it's important at this stage to think that the main Metrica is not finished. We don't think this is the final product. We intend to add to this, refine this, improve things from feedback that we get, from new use cases people suggest to us, new data sources can be added. Basically, it's a foundation that we can build on. We can add to. And in terms of the data sources we use, we have a methodology that we've published, which is how we evaluate RBLs in isolation, but also in concert. So is an RBL adding unique entries for us, or is it actually just overlapping largely with other sources we have? Should we keep on using it, or is it actually causing more trouble than it's adding value?

And for the future, we have all the research that we're doing within the SSR team. So we're looking at things like the uptime of abuse. How

long is harmful content available for? We're looking at mitigation, different mitigation strategies or methods, and things like is the report based on a compromised domain, or is it actually a malicious registration? Because all these bits of information are useful. They're things that we can add into the system to enrich the data, to enrich the information that's available to the community.

So the process will be, once the research is at a stage where we're comfortable that it's robust, we can build a prototype, and then that can be productionized and pushed in to Domain Metrica in a future module. I think that's the last slide. We go to Q&A.

MUKESH CHULANI: Yeah. We're open for questions. Siôn , there's one question on the, where did that go? There's one question in the chat. Will it be possible to track trends and movements of abuse types across registrars over time and correlate it with discounting promotion activity?

SIÔN LLOYD: So we can certainly give the data that plots those trends as we see it in our data. So you can put the particular time scale that's of interest and see the data as we see it. We won't have the external information about what factors may be driving changes of behavior. So we won't be able to mark when promotions occurred or when prices changed or something.

MUKESH CHULANI: There's another question from Bruce Tonkin. I gather the data is gTLD only at the moment. Do you have any plans to add ccTLDs, those that were on the DAR program, for instance?

SIÔN LLOYD: Absolutely. So module one, we concentrated on the gTLD world. We've been here speaking to the ccTLDs yesterday. I was presenting to them. We're basically asking them how they want to be brought on board. What makes most sense that we can provide the most useful data both for them and the community at large?

MERT SAKA: There was one question I don't know if we were answered from Torsten Krause online. He was asking will, sorry just a second, will also information on CSAM as a kind of DNS abuse available on Metrica?

SIÔN LLOYD: I answered it.

MERT SAKA: Okay. Sorry. I thought you missed it.

MUKESH CHULANI: So let's go back to the room. We see hands raised. We're going to have Reg start off. Reg, go ahead.

REG LEVY: Thanks. Thank you very much for the beta version of it. It was really nice to be able to see it before it went semi-public to all registries and registrars. Thank you also for hearing our feedback. I look forward to spam being removed entirely before it is made completely public. Will there be future beta tests so that when, for example, you add spam back in, beta testers will be able to see it first, then it is rolled out to registrars and registries, and then it is rolled out to the public?

SAMANEH TAJALIZADEHKHOOB: Actually, that's a very good suggestion. We didn't think about this because this was the first time we launched it. We've thought of pilot, but the pilot came out to be very useful for us. And if it was useful for the pilot users, then why not? We will do the same.

REG LEVY: Thank you. I would strongly encourage that you continue the same model that the original rollout followed for any updates to the system. The other thing that I wanted to ask is, you said that data is shared to the extent that the licenses that you have with the retail block lists allow you. I would ask that you only use retail block lists that you can share data from, because it is not fair to tell us that we have, for example, three abusive domains and not tell us what those domains are. I'm not saying that you need to give us evidence. We have many different feeds. We're doing our own systems. But, for example, when I logged in during the pilot session, I saw some domains that I wanted to resolve the issues on, and I was not able to view what those domains were. So if you can't

share that data, it's not fair to publish that data because we literally can't do anything about it.

SAMANEH TAJALIZADEHKHOOB: Thanks for the feedback, Reg. This is a kind of, I'm sure you have been involved in the discussions in the community within where we found a trade-off in where DAR was previously where we were only sharing aggregates and it was from the request from the community where they asked for as much as data that can be shared. So transparency. This is what this is basically the design principle of this platform. We try to be as much as transparent as we can, but we cannot also disrupt businesses of these feeds. The assumption is, and this is the rule that we go to, is that we will use feeds that are accessible as much access, so public and accessible by the community, the wider community of ICANN. This is our principle, but also we will not compromise quality over this. We want to use high quality feeds. Yeah, so feedback taken, we have to see how we can work on that. Is there anything you want to add?

REG LEVY: So to the extent that there are domains that you can't tell us what the domains are, I don't see a current way of seeing what feeds those domains are reported on. So if that information is so-called public to me to access and pay for myself, I need to know where to look for it. And once again, this is going to be used by the community to judge how well registries and registrars are doing with regard to DNS abuse. I support this. This is great. I love it. But please be fair to us because we're going

to be taking the hits regardless of what happens. And so providing us with as much data as possible so that we can actually do what we do is going to be extremely helpful. Thank you.

SAMANEH TAJALIZADEHKHOOB: Thank you, Reg. A valid point. And by the way, you can see the name of the feeds in the dashboard. It shows you which feed is showing your domain. It names the RBL provider.

REG LEVY: But only when it names the domain.

SAMANEH TAJALIZADEHKHOOB: Ah, you mean the aggregates?

REG LEVY: So again, if I see that there are three abusive domains and I click through to see what those domains are and then it says, actually, we can't show you those domains, it doesn't then tell me. But if you subscribed to Retail Block List A, you would get those domains. What it does do is you have three abusive domain. You have three abusive domains. They are exampleone.com, exampletwo.com, examplethree.com, exampleone.com comes from RBL A, exampletwo.com comes from RBL B. Thank you for hearing it. Thank you for this session.

MUKESH CHULANI: Neha, your hand was raised briefly. I think you put it down intentional.

NEHA: Reg actually covered one of the questions that I wanted to ask.

MUKESH CHULANI: Okay, that's fine. Thanks.

NEHA: Sorry, but I also wanted to ask if you can briefly touch upon the popularity of domain. I'm sorry if this was covered before and I missed it. And how exactly that is measured.

SAMANEH TAJALIZADEHKHOOB: Yeah, so we have to, actually a Tronco is a popularity list that is developed by academics and Maciej and I are two of the authors of the list. It was initially developed as to be a list so doesn't have the downfalls of the existing lists, popular lists such as Cisco, Umbrella, Alexa, One Million and others. It's it has peer reviewed method and it's been picked up by a lot of industries for white listing and other kind of work. And it's so far the most credible popular ranking out there. What it does is it aggregates the basically different metrics that indicate how popular a domain is on that. And by that, I mean DNS related metrics or network related metrics.

Why we included it is because the research shows, and there has been a body of research on this that there is a relation between how popular

a domain is and it's becoming a target for DNS abuse. So we provided that data so that the community can download it and do their own correlation analysis on it. And it's free of free of you. So if you go to the Tronco website, you can download it for free on a daily basis.

MUKESH CHULANI: Samaneh, one more question, I'll take one final question on this. Oh, sorry. Go ahead, please. And we'll manage time after.

ROB GOLDING: Rob Golding, speaking as a registrant, but also working with registrars and the data sources for this concern me. You've obviously going to be getting fed in details from ICANN compliance and you're taking in details from RBLs. In my experience, the bulk of malware distribution takedowns and notices don't go anywhere near an RBL. They're dealt with by the host long before an RBL ever sees them.

Similarly, some registrars are very, very good at dealing with potential trademark issues, potential abuse of their domains, potential malicious actors simply by reporting it to them. Much more bottom up solution than top down solution. None of that data will end up in here unless you are specifically asking registrars to also supply feeds. This then has the potential to become for the community sort of beauty contest of the people who aren't proactive, who aren't dealing with it quicker than you would see it. And I don't think that's necessarily fair or appropriate.

SAMANEH TAJALIZADEHKHOOB: Thank you for the question. This is a point that we are very well aware of. And that is why in the beginning I set the scope, we set the scope that this is not meant to be a place where you provide information on whether a domain is not into that. We are not going into that scope. This is a place where we are measuring what we see in RBLs. We are saying these are the-- where we see from this perspective abuse are concentrated. We are well aware that mitigation sometimes I've heard before that. We don't have that view. There's no way that we can get from external data that view. That's the data that either registrars have themselves or they are working with third party companies that have it.

In the future, we will design more metrics that allows us to have more robust measure than existing ones on uptimes. The reason why we don't have those metrics yet in, although they exist, is because we do not want to base those on RBLs. We want to be more precise. We want to exactly take what you said into account. So stay tuned. That will come. And I mean, discussion with you guys will be very useful for us because it will give us insight on what you see and we don't see. So let's keep these discussions going back and forth.

DIRK KRISCHENOWSKI: Dirk Krischenowski from Berlin. I just want to echo some talks I had with GAC members about this and other things, and they said it's really good that ICANN is doing a lot in this direction. But in our country, they said the ccTLD is the most important thing because there are 50, 80, whatever percent of all cases are in our country. It's nice to have the

gTLD data, but it doesn't help a lot because the problem is the ccTLD and we don't have that data here in that Metrica platform.

SIÔN LLOYD: Thank you. And that's something we're aware of. But we're conscious of when we bring ccTLD data into the system, we need a level playing field with the gTLDs that we're also reporting on. So, for example, the size of the zone file we don't have because we don't have access to the zone. So we need to work out how we can do it. But in a way that's fair, that the data is comparable across all of the system. But it is going to come.

SAMANEH TAJALIZADEHKHOOB: It has been designed on purpose like that. We want to get consent, like Siôn said, from ccTLDs who are participating in DAR, that they can be automatically migrated to the system. And also, even though there are ways that we can have estimate zone sizes from external sources, we want to have consent from the ccTLDs that we can use the sizes. So we want to be careful and precise. It will come. So stay tuned for a webinar that is happening soon with ccTLDs to discuss this exactly.

MUKESH CHULANI: Samaneh, can we do one final question before we move to INFERMAL? One final one on the chat. So RBLs make money for these lists. If a registrar contacts them and we provide proof that they're not taking action to remove it, the domain, once deactivated, will there be any repercussions for that RBL?

SAMANEH TAJALIZADEHKHOOB: We do not speak for what RBLs provider do. What we do is we just make sure the feeds we get meet the standard that we evaluate the feeds. The rest of the questions to what feed provider does, I think, can be directed to them. They are here with us in the room today. We have, I think, RBL Spam House in. And so further questions can be discussed with them directly. Thank you.

DEAN MARKS: Super quick question here. Thank you so much for all the work you've done on this. Dean Marks, with COA. I just had a very quick question about the test group. Is there a list of who those test users are? So we know what parts of the community are in the test group.

SAMANEH TAJALIZADEHKHOOB: Very good question. Thank you, Dean. We actually haven't published that, but we can and we will.

MUKESH CHULANI: I haven't grown eyes at the back of my head yet. So please, you can.

SAMANEH TAJALIZADEHKHOOB: We are not taking any questions.

VOLKER GREIMANN: Yes, hello. Good work there so far. Very interesting. And this may be outside of the mission of ICANN, but still relevant to the topic. Websites and domain names that have this kind of abuse, they certainly need a delivery mechanism. And many studies have now shown that 90% of all scams, for example, originate on Facebook and related platforms.

Is there a view of looking at those delivery mechanisms, i.e. how people are directed to those domain names and making available some data on that so that we can maybe work with other parts of the community or other communities even in addressing what leads to DNS abuse, i.e. how people get to the domain names? Because that might be a very, very important stepping stone in eliminating DNS abuse as well. If nobody finds those domain names anymore, then DNS abuse will go away as well.

MUKESH CHULANI: Volker, we can take that question later on, I think, in the interest of time, we'll move on to informal. But the question is noted. Thanks, Volker.

SAMANEH TAJALIZADEHKHOOB: So for all the questions and feel free to find us after this room for more questions. Now we are moving to the next slide, please. So I will do a very quick intromagic.

So the second part of this session is focusing on the study that ICANN funded and Kor Labs conducted. This is research conducted by Professor Maciej Korzynski and the team called Inferential Analysis of Maliciously Registered Domains. The study focused on finding factors

and features, proactive and reactive features that contribute to attacker's preferences across domains, TLDs and registrars. Maciej, the floor is yours.

MACIEJ KORCZYNSKI: Next slide, please. So thanks Samaneh for the introduction. So my name is Maciej Korczynski. I'm a professor in Grenoble Alps University. And today I'm going to discuss the results of the INFERMAL project that was conducted by Kor Labs and Grenoble Alps University. And in particular, the entire team, Yevhenia Nosyk, myself, Sourena Maroofi, Jan Bayer, Zul Odgerel and Andrzej Duda. Next one, please.

So to motivate this work. So cybercriminals exploit domain names and DNS system for launching phishing attacks, spam campaigns, malware delivery for ensuring reliable communication between command and control servers and compromised hosts within botnets. And therefore, the attackers need to constantly register new domains. And previous studies showed high concentrations of abuse within certain registrars, within certain TLDs, pointing anecdotally that the price is a main driving factor of abuse.

That being said, there is no comprehensive research that would statistically analyze this phenomenon by including all or a very comprehensive list of potential features that attackers might be attracted to and the factors influencing malicious registrations. Next one, please.

In this project, we bridged this research gap by investigating domain name abuse. What is super important, we're doing it from the attacker's perspective. We identify factors driving malicious domain registrations using GLM modelling. Next one, please.

The scope of this work is on maliciously registered phishing domains. Why not other forms of DNS abuse? For malware delivery, we see that the great majority of URLs do not contain maliciously registered domain names. For spam, some subset of spam will not fall under the definition of DNS abuse. Finally, phishing domains, there are a lot of them, there are a lot of phishing attacks. Plus, they come with evidence based on which registries/registrars can actually take action.

In our study, we collected data on 73 features that we divided into three groups of latent factors. Registration attributes, proactive verification and reactive security practices. Using GLM regression analysis, we targeted two questions, sub questions. We studied the relationship between features and the concentration of phishing domains at the registrar TLD level. Why registrar TLD level? Because some of the attributes, some of the practices might come directly from registrar, some directly from registry. Think about uptimes. Technically, both registry operators and registrars can take action at the DNS level to mitigate abuse. The second research question was around, are features favored by attackers alone, or by both attackers and legitimate users? Next one, please.

To answer those questions, we used the following data sets. Phishing block lists from anti-phishing working group, PhishTank, OpenPhish.

We also sampled benign domain names very carefully using stratified sampling and a lot of different methods to make data representative. We used open zone files from CZDS, from a certain ccTLDs that provide open zone files. Also, we collected domains from Google City Logs.

As for features, one of the main data source that we used is TLD list. This is a service that on a daily basis, three times per day, updates and provides the data per registrar, per TLD, on things like domain registration, retail pricing, discounts, free features that are coming with the domain registration, and so on. A large part of the data we collected manually, things like free unrestricted API access, to for example, create sub-user accounts, APIs to register domain names automatically, different restrictions. Also, we performed active measurements of uptimes. Next one, please.

Here, very briefly, our input data was 500,000 block-listed URLs collected between August 2023 and January 2024. I'll just summarize a little bit of the slide for the sake of time. There are two main principles. One is, the method is much more complex, but the main assumption to get maliciously registered domain names was if they are taken down at the DNS level, if they are delisted from the zone file or removed from the namespace within one month, then we were considering them as maliciously registered.

The second very important point is for each maliciously registered domain name, we were checking the creation date and the registrar, so that we could check in the TLD list what were the registration attributes

offered by that registrar on that given day with that particular TLD. Next one, please.

I'm not going to go through all 73 features, but just a couple that turned out to be important or potentially important in the statistical analysis. The first group, registration attributes, we checked the free APIs, free bulk search, available payment methods, including cryptocurrencies, digital wallets, retail pricing, discounts, also pricing terms, which might apply only to a limited number of domains or new customers, and also free services that are coming with the registration. Next one, please.

The second group of selected features are proactive verifications. We manually created the registrant accounts with different registrars, checked the operational validation of email and phone numbers. We also defined three labels, one random string, and second and third containing keywords, Facebook and Office 365. We selected those two, Office and Facebook, because at that time those were the most commonly targeted brands in our phishing data sets.

For each registrar-TLD pair, we created a domain name and we proceeded to the payment prompt without actually purchasing them to see if there are any registration warnings or if we are prevented from registering them because of specific patterns in the domain names, in the labels. We also collected the data on registration restrictions, like local presence required, ID required, or some TLDs restrict registrants to be for example certain professionals. Next one, please.

Finally, we also measured reactive security practices, because all the time, keep in mind, we were doing it from the attacker's perspective.

The hypothesis here was that perhaps reactive security practices of registries or registrars might influence the choice of TLD and registrar to be abused by the attackers. I will explain it on the example later. We performed WHOIS and DNS measurements, and we were measuring uptimes, meaning the time between block listing and when the domain name was taken down at the DNS level.

For one month, there were more granular scans at the beginning of the scan, as our data shows that they are mitigated generally quite fast. Then after 48 hours, we scan every 12 hours. For a subset of maliciously registered domain names, we also send notifications. Next one, please.

Before diving into regression analysis, I will just show a little bit descriptive analysis of the data of the maliciously registered domain names. Here you see two figures on the left showing pricing. On the right, discounts. On the y-axis, you can see cumulative distribution function. For maliciously registered domain names, we've seen that price ranges from less than $1 to $69. When we take a look at the CDF function, nearly 50% of the domains cost $2 or less. Next one, please.

That being said, we also observe a long tail there. There are, of course, examples of expensive domain names. You have a couple of examples here, including a screenshot of one of those domains. Here the hypothesis simply is that the attackers might have a good supply of stolen credit card numbers, access to stolen, for example, Bitcoin or other digital wallets. They care more about having deceptive keyword in a domain name that might be more expensive, and they might care less about the price. Next one, please.

We also observe that registration prices might be subject to various terms. For example, out of around 15,000 domain names, more than 4,400 domain names had reduced registration price that were valid for one domain only. Also, some discounts were to new customers only. Next one, please.

This is also quite interesting. We analyzed a free and unrestricted API. It was available and offered by 16 out of 31 registrars. Four registrars also permitted an automated account creation, in particular, for example, as a sub-account under the existing API user, which also might be quite appealing from the attacker's perspective, although it is a completely legitimate feature. And then, for example, if we take a look at the three features in the figure and the number of the domain names on the x-axis, we could, for example, take a look at a free API and free bulk search, a significant number of maliciously registered domain names. At the time of registration with the specific registrar and TLD, this option was actually available. Next one, please.

We observed 13 payment methods across analyzed registrars. Also, we observed that all the emails are verified by registrars. That being said, we observed a variety of ways of implementing it, and only six of them validated operationally for number. Also, two registrars blocked trademark domains from being added to the cart. Next one, please.

This is also quite interesting, because here we are comparing uptimes for domains for which we performed notifications and for which we didn't. On the x-axis, you see uptime in seconds. On the y-axis, you have a cumulative distribution function. Both at the domain name level and

at the registrar level, generally, the trends are very similar. How to interpret those results? Most probably, the majority of registrars are already consuming those feeds one way or another and take action on block-listed domain names that are appearing on the block lists. Next one, please. Now, let's dive a little bit more into the...

MUKESH CHULANI: Just time check. You have nine minutes before the session ends.

MACIEJ KORCZYNSKI: Yeah, I can say straight away I will not make it. Okay. Here, we analyzed the relationship between collected features and the concentration of phishing domains at the registrar TLD level. I will start with the initial retail registration price. We see weak positive correlation. Average results show that if we decrease the price by $1, the expected increase in the abuse concentrations, in abuse incidents, should be by 6.6%. By keeping all the other features constant by the registrar TLD and so on.

That being said, those also results need to be interpreted correctly because there is no per se linear correlation there. For example, let's imagine increasing the price by $10. There will be much higher correlation if we increase by $10 the price from $1 to $11. The correlation and expected decrease in the abusive domains will be much higher in comparison to when we increase the price from $100 to $110. Once more, $10. I think it's very intuitive and it's also how this should be interpreted. Of course, it sounds like retail or initial retail pricing, there is a weak correlation, but it does not mean that the attackers do

not prefer pricing. Because we, also, in our analysis, we included retail registration discounts in principle that basically contribute to the final price that registrants, regardless if legitimate or malicious, they pay.

Here, the results suggest that there is much stronger correlation and that the attackers, phishers, value registration discounts and exploit them. If we provide a discount and it results in a decrease in the final price by $1, we should expect around a 50% increase in abuse rates per registered TLD. Both of them contribute to overall pricing. When we take a look at cryptocurrency payments, imagine that some, let's say, a registrar does not provide cryptocurrency payments and then they decide to introduce them, then the expected increase in abuse counts should be by 30%.

What is also very interesting, where we observed the strongest correlation between the driver and abuse counts is the access to API to register domains or to create sub-accounts under the user. Here, when we enable such registration, from such feature, from the attacker's perspective, they would likely exploit it and we should observe a 400% increase in the abuse counts. What is also important that in our analysis, we did a lot of feature engineering here. And some things, for example, cryptocurrencies or as you can see, APIs, some features were aggregated into a concept. For example, we were not verifying Bitcoin and very particular payment providers, but aggregating them to cryptocurrencies or digital wallets. Next one, please.

Now, free services are very valuable also for the attackers, which makes absolutely sense. It saves time for the attackers. They register a domain

name, they deploy a website very fast, they have available DNS service. There is no problem. And here, we see increase in abuse rates as well. Presence of restrictions.

Here, we had a challenge because of the variance in the data. The variance was very low and we had to make a trade-off and we aggregated all the restrictions that we had into a concept presence of restrictions. This is definitely the limitation of this work because we could not analyze each restriction independently and it has negative correlation and if some restrictions are present, then we should observe a decrease in the abuse counts by 63%. Same negative correlation with operational validation of email form present.

What is quite interesting also is that uptimes do not correlate with abuse counts. The idea here, I mean, intuitively, if I'm the attacker, I go to certain registrar and try to abuse the TLD and if my domain names are taken down very fast, next time I would register with some other registrar potentially in TLD. Our hypothesis proved to be wrong and plausible explanation for this would be that there is a significant time between domain name registration, launching the attack, detecting abuse, block listing URL, sending notification, taking action at the DNS level. It is enough time for the attackers from the economical perspective to get enough credentials so that they simply do not care if someone takes down fast or not. Next one, please. Next one, please.

Okay. So, discussion and conclusions. In this research, we analyzed key factors driving malicious domain name registrations from the attacker's perspective. Importantly, the study did not include instances

of benign domains that have been compromised. We made a lot of effort to be sure because that was super important from the statistical perspective.

MUKESH CHULANI: I'm really, really sorry.

MACIEJ KORCZYNSKI: Oh, no, I need to do, I need to say it.

MUKESH CHULANI: Give me two seconds and then you can continue. We're going to get cut off. We'll arrange a webinar before the end of the year to continue with any post questions. There are questions we can't get to them. There must be questions in your head. We won't be able to get to them. I'll turn back to Maciej and we'll try to arrange something again.

MACIEJ KORCZYNSKI: Okay. Yeah, that one is actually one of the most important slides, so we couldn't really skip it. I'll be available for questions, of course. No, no, no. Excuse me. We need to go.

MUKESH CHULANI: You can have a few.

MACIEJ KORCZYNSKI: Few minutes more. Okay. Thank you. So, also, results really need to be interpreted with caution and may or may not be generalized into actions by registrars or TLDs. This is because we took really the attacker's perspective. I can give you some examples. The results clearly show that if you disallow, for example, cryptocurrencies, the expected decrease is by 30%. Price counts here. But, for example, the attackers value DNS and free DNS free hosting service. But it does not make any sense to impose on anyone to take this service out from the registrars because it would be simply insane. So, my point here is this is from the attacker's perspective and we need to very carefully think how to translate some of those into some actions.

Some factors are combined and involve a variety of features like registration restrictions. Sometimes it's a limitation, sometimes not necessarily. What is important also is that attractiveness to the attackers likely results from a combination of factors. Here I would like to give an example. If I were an attacker with Bitcoin available, I want to streamline my operations, I would go with a registrar that provides, let's say, an API to create automatically subaccounts and, of course, automatically registering domain names, free DNS, of course, free hosting and, let's say, also promotions per user, which is not a big deal because I'm creating subaccounts that are eligible for those discounts and I pay with Bitcoins.

What I'm trying to say here is the attractiveness to the attackers really is a result from the combination of different factors. Now, my final thought. It is really essential to consider economic implications of all of this. We're showing here just part of the picture. From the attacker

perspective, we show we do this, this is expected outcome. But out of the scope of this study is what are the economic implications, how much it will cost to develop policies to put technical solutions.

We tried to investigate this a little bit, but in practice, the impact on legitimate users is unknown. We may speculate that maybe if we disallow crypto, maybe it will not affect that many legitimate users. But it's speculation. We need to also make risk analysis because the attackers will likely respond, not will likely, they will respond to those adjustments and they will bypass them. We need to constantly increase barriers for abuse.

But also, we need to think long-term what the attackers could do. I have a couple of examples but I guess for coffee break. Thank you so much. Last slide also. Thank you. No, no. Acknowledgements are very important, actually. Many thanks to ICANN for funding this research. Many thanks to Samaneh and Carlos for a great, great discussions. And also, data set providers. Without access to those data sets, this research wouldn't be possible. Thank you so much.

MUKESH CHULANI: Thank you all for being good sports. We'll try to arrange a webinar for the Q&A portion. Thank you.

[END OF TRANSCRIPTION]