ICANN83 | PF – At-Large Policy Session with OCTO
Thursday, June 12, 2025 – 09:00 to 10:15 CEST

YESIM SAGLAM: Hello and welcome to the At-Large Policy Session with OCTO. My name is Yesim Saglam and I'm the remote participation manager for this session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior and the ICANN Community Anti-Harassment Policy. During this session, questions or comments will only be read aloud if submitted within the Q&A pod.

Interpretation for this session will include English, French and Spanish. If you would like to speak during this session, please raise your hand in Zoom and when called upon, virtual participants will be given permission to unmute in Zoom. On-site participants will use a physical microphone to speak. Please state your name for the record, the language you will speak if speaking a language other than English, and please speak at a reasonable pace. Thank you all for joining and I will now hand the floor over to Jonathan Zuck, ALAC Chair. Thank you.

JONATHAN ZUCK: Thanks, Yesim, and thanks folks for coming on board here. This is sort of our DNS abuse day. We're having a lot of conversations about it today and we've been meaning to connect with OCTO to

understand the research in which you're engaged and plans around DNS abuse and how it's measured and tracked, etc. Perhaps look at a demonstration of Domain Metrica, etc. So, I'm glad we're finally putting this together. So, I think without further ado, I would go ahead and get a little introduction of OCTO from Samaneh. So, take it away.

SAMANEH TAJALIZADEHKHOOB: Thanks, Jonathan. First of all, thank you for having us and inviting us for today's session. We are happy to be here and also talk about our work. And with that I will start today's presentation. So, today it will go as follows. I will give a brief introduction of what are the general themes that we from SSR OCTO work on and what is our team and what we do. Then I hand into Sion to talk about Metrica and then Carlos and Sam from my team will present a brief overview of some of the other research we do inside SSR OCTO. And lastly, I'll close the session with talking about the INFERMAL report.

Next slide please. Yes, so let's start from the very definition that we in the community have been talking about for years. When we think about DNS security, DNS abuse or DNS insecurity, we talk about domain and DNS infrastructure and its associated components and processes that are not adequately secured. The reason why I'm highlighting these words is these are the points that we focus on later on during the presentation.

When these processes and components are not adequately secured, it creates opportunities for malicious actors to exploit them either via vulnerabilities or other sources that are there including the weakest links to engage in various forms of DNS abuse such as phishing malware, command and control, spam as a delivery mechanism and many others that are outside of the current remit that we defined inside the ICANN community.

Next slide please. But then we think why does DNS abuse, specifically DNS abuse exist? There has been work, a lot of work on this from computer scientists that mapped computer science as a discipline to economics, basic economics. And the idea is that the security failures of systems is caused at least as often by bad or misaligned incentives as bad design. So basically, systems are particularly prone to failure when the person guarding them is not the person who suffers when they fail.

Next slide, please. Why this is relevant is because we often, in our community, talk about these things but sometimes we do not connect them with each other. When we talk about guarding systems we talk about proactive security measures. So, those who technically make a system secure in order to basically prevent from malicious use of it later in the future.

When we talk about those who suffer when systems fail, we talk about reactive security measures. So basically, first failures happen and there are victims of a failure and then those who guard the system might take actions which we call them reactive security.

But the two are sort of related but also sort of unrelated because one is preventive, one is reactive. And this is the area that we work on basically and the idea is that these two should be aligned with each other. So basically, proactive measures, the guardians of the system, their incentives should be aligned with those who are the victims of the system, which we know in our space is not the case.

Next slide, please. Then we think, okay, if the incentives need to be aligned, how do we do that? This is one of the suggested ways to do based on literature on economics which exists for a long time. It is designing high-quality robust metrics. This is regardless of the field. We are now talking about DNS abuse but this is regardless of the field. What are high quality metrics? Metrics influence decision making by quantifying security performance, allocating resources and impacting priorities. Poorly designed metrics can create reverse incentives leading to unintended consequences and short-term progress. What do I mean by that?

Next slide please. Can you click next so that I see next? So here on the left what we see is basically that both of the plots are showing the concentrations of domains that are reported as DNS abuse. The plot on the left is showing the absolute count, so basically the number of domains that are reported on these lists. Sorry. The plot on the left is showing the counts. The plot on the right is showing the percentage. Both of them are showing concentrations but in very two different ways.

The plot on the right is taking the size of the TLDs into account. If you look at the blue dot which is an example of TLD basically, I call it .CC, just not to point to any TLD. As you see, once you take a factor into account, once you normalize a metric the position of this dot changes a lot. It is this kind of the same metric, the same idea but shown it's been calculated in two different ways. This shows and looking at these two different plots, one actor or a decision maker can have very different interpretations showing how important it is to have a good metric which is relevant to the decision you want to make.

Next slide please. Then we think what makes metrics better? High quality data, data with fewer false positives, more precise definitions of what is measured, for example spam as unsolicited email versus spam as a delivery mechanism for malware. These two definitions creates very two different metrics. Incorporating measurement errors, acknowledging limitations of the metrics that are used.

We often hear in our community that there are publications on the trend of phishing or the trend of DNS abuse but there's little information in the reports about how is this measured and what is the timeline, what is the level of measurement, what is the level that is count as a malicious activity. Because each of these definitions can really differ, can create various set of plots. All of them are true but they are just different from each other.

Next slide please. With having that said, who we are. We are the SSR research team, Security, Stability and Resiliency research. We sit inside ICANN's office of CTO, OCTO, together with the other two teams, Research, and Technical Engagement. I lead this team, and my team consists of Samuel Cheadle, Sion Lloyd and Carlos Ganan who you will hear from them today.

Next slide please. Basically, what we are, we are a team that mostly do research and development on internet identifiers. The core idea of our work is that we look to use data and scientific methods to answer questions that are of interest to our community, talking about the ICANN community but also scientific community. Our goal is to increase reliability of the metrics that we measure, whether they are security and security abuse. DNS abuse is one of the topics we work on. Our focus is on areas where existing methods and metrics need improvement, or we work on developing new methods from scratch.

Next slide please. The reason why I said what I said up to here was to introduce ourselves and basically to give you before introducing ourselves an overview of why the things we work on are important and then actually go into the detail of what we do. From this point onwards I will walk you through what we do in more detail.

Let's talk about what are the common pitfalls in current DNS abuse metrics, the existing ones. We have inconsistent abuse definitions, which is a problem that has been discussed in the ICANN community for years. There are instances where the edge cases

and false positives are overlooked for their use case, because in some use cases they are not that important, but in some use cases they are. The context is ignored sometimes when there are discussions in our community about DNS abuse, because the context matters a lot, especially if we are talking about, depending on the level of analysis, the counts and how we look at these counts matter a lot and they are different from each other.

Even though when we talk about TLDs, registrars and hosting providers, we talk about it sounds like they are three different entities, economic entities, but we know for a fact by looking at their markets is that they are operating in very different markets, simply markets. So, it matters what kind of metrics we define for them to capture the incentives within that market. Some of the most important pitfalls of the common abuse metrics, which I want to focus today and actually ask the community to pay attention to, is a lack of transparency in data sources and methodology of the metrics published and the so-called RBL paradox and the temporal Simpson's paradox, which I will talk about more in this presentation.

Next slide please. What do I mean by RBL paradox? First of all, RBL stands for Reputation Block Lists. If you haven't heard of this abbreviation or this word, these are kinds of lists that basically list domains that are important to be used in a different source of malicious activity. When we talk about RBLs, we talk about activity

that is within the scope of our work, which is phishing, malware and command and control.

Now going back RBL paradox, many eyes, no vision. So, what we say is that most of these lists basically identify threats, different sets of malicious threats. What we saw in our own analysis, if I'm not mistaken, OCTO 37 document and the scientific article that we published in anti-phishing working group conference, is that the overlap is very limited of these lists. The more threat feeds we use to gain clarity, the more disagreement we expose, because the lack of overlap means that there is disagreement or that we need more coverage. Relying on a single RBL creates blind spots. Using multiple lists reveals contradictions.

RBLs reflect different data sources. They have different update rates. Their threat definition is different. All in all, the more RBLs we use, the more inconsistent-- With each of them, we need to be careful. And we also know that each of them covers one area of this landscape, of the threat landscape. This goes to the example of blind man and the elephant, which each blind man touches one part of the elephant, the landscape, but none see the whole.

Next slide, please. The previous one was one side of the problem, which is we each by using a number of reputation block lists, we are seeing some part of the landscape. But then there is the time aspect of things, the time window distorts metrics. When we are looking at abuse counts within these lists, we know that they are heavily related on the time window that we select them, that we

count them in, whether they are weekly, monthly, whatever time window that we are counting them in.

Short windows miss persistence. So, a narrow time window may miss long-lived abuse domains or delayed mitigation, underestimating the impact. But again, when we choose long windows, we basically smooth over spikes and responsiveness, which in turn hides basically short-term abuse or those fast remediations. The point is different parties use different time windows when they are making abuse reports or plots or trends. The choice of the time window can manipulate on intense or unintentionally skewed narrative, turning the data into some opposite story.

Point is, depending on what the analyst or the reporter chooses to present the data in, the metric can vary a lot. Again, a clear example, phishing went up by 10% from last July. This doesn't say anything. This doesn't say anything. Nobody can react on this based on just a sentence because it depends on the time window that the data is aggregated in. It depends on the number of lists that are used.

Now I'm going to show more in the next slide, please. Another aspect of time here in this plot is basically top 10 ccTLDs are ranked by the amount of malicious domains in them, and then in each column they are aggregated monthly, quarterly, semesterly, and yearly. The ranks are from 1 to 10. They are the same. In horizontal, you see TLD 1 to 10. So, you see that when we look at

monthly versus all the other ones, the rank changes, indicating that shorter time aggregation can create different metrics than quarterly, semesterly, and yearly. Then depending on how is a metric reported can create very different results.

Next slide, please. This is another aspect of time. Each of these bars are one TLD, and what we see is that-- The three colors show basically the percentage of active, inactive domains and those that do not have addresses. And these are only for newly observed domains across different TLDs. So, you see how newly observed domains across different TLDs, they have very different percentages when it comes to the amount of domains that are active. By active, inactive, we have a definition which goes beyond this presentation, but basically those that serve content and that can be queried active on internet. Again, depending on what we choose to show you, our interpretation of this metric could be very different, so you would see something very different.

Next slide, please. It is very important what we look at the data and what we define the metrics. This is what we focus on inside the SSR research team. We generally are interested in security and abuse of identifiers, internet identifiers, which means names and numbers. We look at data and metrics. Next slide, please. We try to find areas where if they exist, how can we make them better? If they do not exist, how can we create metrics that create incentives for more security? Some of the instances of our work, you see in pink, we work on identifying bad registrations. We work on detecting parked domains, classifiers to classify threat types. We

look on abuse concentrations and counts, and we look at uptime of domains.

Next slide, please. This area that is highlighted in the circle is basically the core idea behind ICANN Domain Metrica platform. It is a platform that will focus on data and metrics and improving them, presenting them, etc. With this, I finish my part, and I would like to hand it to Sion to walk you through the detail of the metric platform. Thank you.

SION LLOYD

Thank you, Samaneh. Hi, there. My name is Sion Lloyd. I'm going to talk about Domain Metrica. I'll do a bit of an introduction, but then I'm actually going to go through a case study, so an example of an investigation that you might do with the platform.

So, as you've already heard, Domain Metrica is a platform that we built in-house, and the idea is that we collect together as much data, metadata, pieces of information about internet identifiers, and we bring those pieces of data together so that we can incorporate them, aggregate them, combine them in different ways, and produce as useful a metric as we can from the data that we start with.

And then producing data is only half of the story. The more interesting, the more useful aspect of it, if you like, is being able to share that data, being able to visualize it and interact with it. So, we have two main ways of getting that. There's the front end, the

web front end that we'll see shortly, and then there's also an API where you can get data programmatically, so you can get the raw data, maybe combine it with a data set that you hold within your organization and do stuff that we couldn't think of or do because we don't have that other data set.

In terms of availability, Domain Metrica is open to anyone with an ICANN account, so the account that you use to register for this meeting will give you access to Domain Metrica. And then further to that, contracted parties, so registries and registrars, they get further insights into their own data, so data on domains within their own remit.

Next slide, please. So, there are other platforms out there that do very similar things, but we're trying to build something that's unique, that fills a niche that we had for ourselves that we wanted, and then it's hopefully useful for other people as well. So, we try and get as much data in there as possible, so we have lots of different reputation lists, and that's our way of trying to deal with that RBL paradox that Samaneh mentioned earlier. The more different data we can get, the more different perspectives we can get, then the higher our confidence is that we see as much of the picture as possible. It doesn't mean we see all of the picture, because nobody does, but we're trying to just at least see as much and be aware of where we may have limitations.

And then we have some other aspects, so we've talked about the way that you can interact with data. We'll see some of that later as

well. And then, as far as possible, we're sharing all of the data that we have. Obviously, in some cases, there are contractual limitations, but as far as we are allowed, we share the data that we're holding. We're transparent about what the data is. And then we're also looking to expand to get new data sets, to get new uses of the data we have, new ways of visualizing it, basically evolving the platform to be as useful and interesting as possible to as many different areas of the community as possible.

Next slide, please. So, in terms of the time paradox, we deal with that basically by allowing the user to select what time range they want to look at. If you want to look at a year's worth of data, that's possible. If you want to look at a month or a week's worth of data, that is also possible. And we will show the data and scale it as the time window changes.

As much as possible, we filter out domains that we don't think should be present. And this is a process that we're going to refine and improve over time. So, currently, we remove domains that we know don't exist. So, we get the RBLs. We also have the gTLD zone files. And so, any entry on an RBL that doesn't have a corresponding entry in a zone file, we remove from our statistics. Because that domain doesn't exist. It may have been suspended already. It may never have existed if it's a false positive on the list. So, we don't count those.

And yeah, let's go to the next slide, please. So, that's a basic introduction. And the request for this session was to have a case

study. So, I've kind of hung that on a recent story, a recent blog article that Contractual Compliance released about the gTLD .top. And .top was on a breach notice that has recently been cured. And there's a blog article there. You can read about it. So, I thought it might be interesting to use Metrica to look at that story to see what data we hold and how you can look at that data and see what it actually might mean. So, I've got a bunch of slides now that are going to visualize like a journey through the front end telling us about how we might move from data point to data point.

Next slide, please. So, this is the screen that you're presented with when you log into Metrica. And there's various things going on here, but mostly what we're showing on this page is headline statistics, if you like. It's a broad range of numbers and a breakdown of those numbers across an unfiltered view of all the data that we're holding within Metrica right now. So, it's the number of domains, the number of zones that those domains belong to, the number of those domains that have reported abuse and so on. And then we have the breakdown by the different abuse types. So, phishing, malware, botnet, command and control.

Next slide, please. And we still have some glitches in the system. So, that number shouldn't read as zero. There's lots of moving parts to get all of the different data sets into the right place. And we have a slight issue with one of them right now. So, yeah. Just only up to that one. Next slide, please. The other thing that we show on the front page is 12-month trends. So, how have these

things evolved over time? What are we seeing as the direction in which these things are headed?

Next slide, please. So, the facility that we're going to use today is the search facility for TLD. We can also search on registrars and we can search for specific domains. The search for registrar gives very similar data to what we're going to see for this TLD journey. The search for domains gives slightly different data because it doesn't make sense. We're not aggregating anything. But we try and show metadata about the entity that you've searched for.

So, that's the top section you see there. It's just a bit of background information, the zone size. If it were a registrar, it would be their portfolio size and so on. And then we see the current reported abuse for that entity. So, again, we have the breakdown by the different abuse types. We're also showing the percentage size there. So, that's the count of unique domains. But normalized to the size of the entity. So, the size of the zone file in this particular case.

Next slide, please. And then we have some trends. And we have the proportions of abuse in each of the different categories that we're measuring. So, these graphs here, these charts, they're interactive. We can click into one. So, if we go to the next slide, please. So, we're going to click into that reported abuse count. So, this is a monthly trend, a 30-day trend showing daily counts of the reported abuse. And how that's evolved over time.

So, there's a number of things we can do now that we're interacting with this chart. For example, if we go to the next slide, we can change the time frame. So, here we've gone back to the beginning of April. And because we've got a larger time window, now instead of displaying daily data, we're displaying weekly data just so that we're not having too many points on the chart. It very quickly becomes confusing. So, this is the median count per week. So, that's the way that we're aggregating the week's data into one point is by taking the median for that seven-day period.

Next slide, please. So, looking at an entity in isolation tells you how it's changed, but it doesn't really tell you too much because other things may be happening across the board. So, one thing that we allow is to add other TLDs into the comparison. So, here we've added a different TLD .pro. And we can see that while .top had a big fall and a gradual decline, over the same time period, .pro seems to be heading in the other direction. And although the absolute counts are less, it's definitely trending upwards rather than down.

Next slide, please. So, one of the problems of comparing different entities is that they're actually a different size. So, the .top zone is a lot larger than .pro. So, what we allow is to, instead of plotting raw counts, absolute counts, we plot the percentages. So, we're still taking the weekly medians here, but now it's normalized to the size of the zone that we're looking at. And this shows us a very different picture. If we're considering concentration of abuse or proportionate amount of abuse, now we see that actually while

.top's been falling, .pro is actually sort of as a proportion of the zone file has actually overtaken. It's actually swapped places, if you like, with .top. And we can add more TLDs into the comparison.

Go to the next slide, please. Here we've added the TLD .win. And this shows some of that kind of short-term activity that we might see. We don't know how this is going to evolve, but there's definitely a very sharp increase to a peak around mid-April, and then it seems to be tailing off, although not as quickly as it increased. So, it may be a single campaign is targeting that TLD, or it may be something more systematic, and we might see something over time that makes it look more worrying. So, this is now just showing three different TLDs and the behaviors of those over time. We're actually just about to introduce a new piece of functionality. So, this is a bit of an exclusive here today.

If we go to the next slide, we have the option to compare with all other TLDs. So, what this does is it takes the aggregate of all the other data that we're not showing individually. So, that's all the TLDs except for top, pro, and win, which are already plotted in this chart. And it's showing how those have evolved over the same time period. And that's the diamond-shaped line you can see at the bottom.

So, this gives you general context of the situation. It's saying, well, what's happening across the board? Is there something strange in the data set that could be causing these effects? Or is it something that's individual to the TLDs that we're looking at here? And you

can see that actually the general picture across the board in the same time frame is pretty flat. There is some variation, but on this scale, it's invisible. It's not that large.

Next slide, please. I think the final thing to say about the front end is that all of the data that's behind these charts is available to download. So, there's a couple of links on those charts. You can get it as CSV or JSON format. And, again, the idea is that you may want to plot it in a different way. You may have data sets of your own that you want to combine this data set with that we can't do. So, you can take the data out of the system, into your own systems, and then use it in any way that you see fit.

Next slide, please. The other way to extract data is via the API. So, if you're happier writing little scripts to interact and extract data, or maybe you want a daily update of the data that's in the system and you want something to run and just extract it into your own data warehouse, then you can use the API for that. And it's got all of the same data that's behind the web front end. It's just available programmatically. And this is just a screenshot from the API documentation that we have. It's all there. You can search for all the same stuff. You can put in time windows, everything that you can do on the front end.

Next slide, please. And this is something like you'll get back. So, it's basically just JSON structured. It's very easy then to import it into Python plotting or something. And just like I say, just generate your own charts, combine it with your own data.

Next slide, please. So, just to close about the main Metrica, it's not finished. What we have today, we're thinking of it as the start of a journey, if you like. We've already received feedback from users. We have our own to do list, our own use cases that we want to introduce into the system. And it's very much we think of it very much as a foundation. We're going to build on it. It's going to evolve over time.

I've actually mentioned this in the chat, but we evaluate the source data that we use to try and make sure that we have the best set of data available. Of course, the counter argument to that is if we're plotting long-term trends, we can't chop and change the input data that frequently. We need to do it quite carefully. We need to note when something changes. So, any change in trends can be attributed to something happening with our data.

And finally, there's just a couple of links there. There's a page that gives more introduction, more information. It's got links to the FAQ, to the API documentation and so on. And then the link to the platform itself. And as I say, if you go there with your ICANN account, you can log in and interact with the system. So, I don't know if we want to take questions on Domain Metrica and now, or if we want to wait till the end of all of the other presentations. I'm happy to do either.

JONATHAN ZUCK Thanks. This is Jonathan Zuck, ALAC chair. Thanks for that presentation. I'm looking here online and around the table to see

if there are questions. I have a brief one for you. You made mention of comparing the domain names that are in the RBLs to the zone files to look at false positives and eliminate those domains. And if they existed historically, they're not eliminated. It's just when you're bringing in new data, I assume.

I guess the broader question that seems to come up with some persistence is the fact that a lot of these attacks take place over a very short period of time. And so, it's got to be difficult to measure whether something's actually been maliciously registered because it looks like it's being parked for most of the time, except for the short time it's being used, etc. And the same thing might be true of a takedown if it was used before as it's still part of the historic data, even if you've eliminated it in the current data. Does that question make sense?

SION LLOYD Yeah. So, well, I'll interpret it and you can tell me if I'm answering the question.

JONATHAN ZUCK Perfect. Thanks.

SION LLOYD So, yes, there's a difference between historic and current views. What we do is every day we look at the situation as we see it that day, and then maybe tomorrow a domain is still on the RBL, but is

now no longer in the zone file. So, now we would count it today, but it will not be counted tomorrow, if that makes sense. And we know it's not a perfect system. We have other ideas and ways which we're going to try and refine the domains that we list, the ones that we believe to be actively abusive at any particular time. So, this is very much like the first step in cleaning those RBLs of domains that we don't believe should be contributing to statistics that day. But they may well have contributed yesterday or for the preceding week or for however long they've coexisted in both places.

JONATHAN ZUCK Thanks. Yeah. One of the things that we've heard a couple of times is that when the study was done about COVID-related domain names that had COVID-related terminology in them, that it looked like the huge majority of them were parked. And the assumption is that they were parked for commercial purposes, like by domainers in case people wanted to buy them. How is that determination made versus parked until the brief instances when I want to use it for malicious purposes?

SION LLOYD Again, it's very hard to have a complete view and one that is across all of time. We can only see what's there currently. So, we look right now and we do the best to give our belief to be the current situation. And realistically, if we look now, it could change in half an hour. And we can't allow for that to any great degree. We can

get better at it, but essentially, there will always be that mismatch between when you measure something and when someone looks at that data, whether it's still relevant or not.

JONATHAN ZUCK: Thanks. I see a hand up from Rowena.

ROWENA SCHOO: Hi. Thanks. It's Rowena Schoo from the NetBeacon Institute here. I just wanted to say thank you so much to Samaneh and Sion for this excellent presentation. It's really interesting work. I really like your new charts that allow the comparison of TLDs. I think they're really interesting. And one of the challenges, and firstly, I totally agree with you normalizing it to percentages so you can compare the kind of rate of abuse and having options to look between those.

I wondered in the interface, if you hover over those comparisons, are you able to also see the number of domains? Because sometimes I find that is also quite relevant. We see a lot of TLDs that have quite low numbers overall, and sometimes it can be difficult to reflect that in the data, especially if they have quite small zones as well. And then the second question, I was wondering, when you compare against all TLDs, is that all gTLDs together or do you make any exclusions for small numbers in terms of abuse or zones? Thanks.

SION LLOYD

Yeah, thank you, Rowena. Yes, so you do get more information on hover. You get the numbers. The other thing that we're looking to introduce is the ability to turn off particular lines in those charts, because if you're only interested in malware, it's quite hard to see them against phishing because they get squashed down onto the axis. And all other TLDs is literally all the other gTLDs that we have in the system. So, there is an aggregate of how they all look over the same time period. And yeah, so if you only had one TLD in the chart, then it would be everything apart from that TLD. If you have three, then it's everything apart from those three.

ROWENA SCHOO

Thanks, Sion. That makes sense. I think just something I see a lot from looking at the underlying data is just how many don't have very high numbers of abuse. And that's always tricky, I think, to deal with in these types of charts. But yeah, I think just to reiterate, I think they're excellent. Thank you so much for this work for the community.

SION LLOYD

Thank you.

JONATHAN ZUCK

Thanks. We've got a little bit of a queue online here. Peter keeps going to sleep in between. So next is Justine Chew.

JUSTINE CHEW

Thanks, Jonathan. This is Justine. Thanks for the presentation, Sion. I just had a couple of fundamental questions. One is you mentioned that your team actually does some data cleaning, if I could put it that way, to remove false positives and things like that. So, do you follow a period, a strict periodic timeline in terms of introducing new data or is that automatic? The reason why I ask is, for example, if I say take a snapshot today with the existing data that you have and then I do it tomorrow, would the database be the same? And the second question is, can you please remind me whether there's going to be initiatives to introduce ccTLD data?

SION LLOYD

Yes. So, to the first point, we re-evaluate. The plan is to do it annually. Domain Metrica hasn't existed for a year yet, but the intention is to do it annually. And any changes we make will be announced and will be done carefully pre-announced so people will be aware that it's happening. We might put a notice up on the system itself to make it clear. We're not going to change the data and have that invisible to the user, if you like.

And in terms of ccTLDs, yes, we're looking at how to introduce them. Some of the data cleaning we do currently relies on having the zone file and at least knowing the size of the zone to be able to normalize it. So, we're just looking at how we can address potentially not having access to the zone file, but still have the TLD presented in a comparable way to the gTLDs that we have.

JONATHAN ZUCK Great. Hadia?

HADIA ELMINIAWI Hi, this is Hadia. Thank you so much for the presentation. And I wanted to ask, how does Domain Metrica compare or relate to DAAR? So, DAAR did include ccTLDs. So basically, if DAAR still exists and the difference between the two platforms and also accessibility. So, who can actually access or use Domain Metrica? So can the community, for example, use Domain Metrica and how? Thank you.

SAMANEH TAJALIZADEHKHOOB Thank you for your question, Hadia. So, DAAR still exists. The platform is working. But we are working towards depreciating it. The idea of Metrica, so Metrica is modular and we intend to add a module in each time period. I'm not specifying them because they are not clear, but the first module was about DNS abuse. But the platform is not necessarily bound to present data on abuse. We are working towards presenting data on all kinds of metadata, on all kinds of things related to domains, IPs and security. So, in that sense, Metrica is way beyond DAAR, and even the first module, which has DNS abuse metrics, is beyond DAAR.

About the ccTLD integration. So, we had, I think, around 40 ccTLDs participating in DAAR. We are intending to add a ccTLD endpoint to Metrica. Sion and the team hold sessions with the ccNSO DNS

Abuse Working Group, to brainstorm with ccTLDs on what they would like to have and how they would like to see it. We took that feedback into house and we are now working towards implementing them. And basically, in the future, we will have ccTLDs being able to log in through their own credentials and seeing their own data in addition to all the other public data.

For the participating ccTLDs. So, we have consent from those who participated in DAAR to automatically move to Metrica. We would like, and we will call for this, for whoever that is interested to join Metrica to come forward and contact us and provide their zones. But ccTLDs can also take part in Metrica if they do not want to provide their zones. We are working towards solutions where we can just get the zone size from the TLD to be able to calculate the metrics where we need the size. So, stay tuned. You will hear more about all these things in the future when we are working towards them. I hope I answered your question.

With that, I would like to ask to resume the presentations in the interest of time so that we, if possible, leave the questions to the end.

SAMUEL CHEADLE

Hello, my name is Sam Cheadle. I'll be presenting a project that's currently underway on Batch Registrations. This is a piece of analysis where we're focused on trying to identify related registrations. There is some conversation and discussions in the community currently around ways to try and identify related

domains at the time of registration. One of the main drivers for this particular piece of work is to take the centralized data source that ICANN has access to and to test the feasibility of identifying groups of related domains, bearing in mind that this centralized data source holds no personal data, so we're not able to tag domains based on the registrant or the account holder.

As we know from a number of studies and past investigations, attackers often register domains in bulk. We've seen many examples of this in the past. We know, for example, that APIs are commonly used. The recent INFERMAL report confirmed this. The bullet point describes this as a hypothesis, but I think it's stronger than that. There's a number of different pieces of evidence. APIs are, of course, used for many legitimate purposes, but they're often used by attackers to register big bulk registrations.

What we're focusing on in this particular study is batch registrations. Those are time-bound groups of domain registrations very similar to bulk registrations, but whereas bulk may cover a slightly longer time period, we're focused here on shorter time-bound batches. One of the key motivations of this work is to try and pick out these groups of domains quickly, enabling more proactive threat mitigation, and exposing these results to the community.

Next slide, please. As I mentioned, this work is focused on a centralized data source with a limited number of registration features. This is akin to a thin registration data set, so we're just

looking at, in this case, registrar, authoritative name servers, and creation dates. So, the basic method is to separate all recent domain registrations based on the registrar and the authoritative name servers. Once we've got those subsets, then we run a clustering algorithm. It's DBSCAN. It's a density-based clustering algorithm. We run that on the timestamps, the raw creation date timestamps.

Of course, this method depends on, to some extent, on the popularity or the frequency of registrations within each registrar name server grouping. There are some cases where there are many, many domains being registered in short spaces of time. Often, those are associated with common default or shared name servers, so those are excluded in this current phase of analysis.

Next slide, please. As I mentioned, centralized ICANN data, we're looking currently at BRDA data. That's bulk registration data access. gTLDs only, no ccTLDs, and for some of the analysis, we are cross-checking with reputation block lists, a number of those that we hold internally, and those are listed in the slide.

Next slide, please. And one more. Thank you. So, just to give a very high-level overview, this is what we see for a one-month period after running this clustering algorithm to try and pick out registration batches. What we have here is, along the y-axis, the 25 registrars with the highest rates of batch registrations. Along the x-axis, we have time. Each of the dots, either blue or red or in between, show the volume of batch registrations. So, the larger the

dot, the more batch registrations per hour. The largest dot is reflecting around 4,000 batch-registered domains per hour, and the smallest, about 10.

Some of the key points to take away from this graph, not all registrars have an equal proportion of maliciously registered domains. You can see that there are some with consistently high abuse rates, shown in red. You can also see some registrars that have large-volume batch registrations but are relatively infrequent, and those may be indicative of campaigns.

Next slide, please. So, just to give a summary of the key points of the analysis that we've run so far, and the main results, so just extending that time window a little bit, we're looking at the first quarter of 2025. What we see, if we look at all the newly registered gTLD domains, we're looking at about 16.5 million domains. If we run this clustering algorithm attempting to pull out batch-registered domains, we're tagging around 4.2 million, so that's just over 25% of registrations. Bear in mind, batches can be any size, so we are including small batches, two, three, four domains, so that's just one point to bear in mind. Often, analysis is looking at bulk registrations and looking at larger batches.

Next slide, please. If we look at abused domains, just those domains popping up in the RBLs, within that three-month period, we see around 1.5 million newly registered abused or malicious domains. Of those, around 770,000 domains we are tagging with this batch clustering algorithm, so that's just over half of them. As

expected, that's a higher proportion than the general frequency of batch registrations. Splitting by threat type, we can see that the spam abuse type has the highest proportion of batch registered. Commander control and malware, the lowest, but they're all above 25%.

Next slide, please. Just to give a brief flavor of the TLD composition. Once again, we're only looking at gTLDs here, no ccTLDs, but if we break down the batches into those that have been tagged by RBL feeds, so malicious batches versus everything else, so we don't have any evidence of threat for these unknown batches, we see this kind of composition. So, these pie charts just showing the top five most frequently observed TLDs. A number of gTLDs that we see fairly high abuse rates for or we have seen in the past also show high batch registration rates.

Next slide, please. Just following up on that point, if we look at the relationship per registrar, so in this plot, this is a scatter plot, each of these dots is showing a single registrar. On the x-axis, we've got this batch registration rate, this metric that we've been looking at in the past slides. On the y-axis, just conventional percentage malicious or percentage of domains within the RBL feeds, and what this slide is highlighting is there is a clear relationship. Those registrars with higher batch registration rates also have higher abuse rates in general. The size of the dots in this graph is showing the volume of batch registrations, the number of domains, so significant positive correlation.

Next slide, please. So, this was a very brief overview. Hopefully, I've been able to highlight the key points. This is an ongoing piece of work, and we're very interested to get feedback on this. We also are currently in the phase of validation, so we're very interested in working with registrars to try and evaluate the reliability of these results of this clustering methodology that I've described, but we're confident that, on the whole, we are picking out valid batches. These domains are related. There may be some edge cases and false positives. That's the issue that we're currently investigating. We do need help from registrars because we're not able to look or access account holder or registrant-level data, so collaboration needed.

Ongoing work. We'd like to broaden this out a little. We don't just want to look at these time-bound batches. We want to look at a more ideally be able to do this on an ongoing basis to be able to pick out these new registrations, these new batches quickly. I'm very happy to answer questions probably offline. I think we're a bit pushed for time at the moment, but please get in touch either via email or after the meeting. Thank you very much, and with that, I'll hand over to Carlos.

CARLOS GANAN: Hey, thank you, Sam. Good morning, everyone. Unfortunately, we won't have time to go through my current research. So, my name is Carlos Ganan, and I'm currently focusing on analyzing domain abuse uptime.

Next slide, please. I'm just going to present the idea to leave room for Samaneh to present the INFERMAL project. As we can see till now, most metrics count how many bad domains exist. My research is different because instead of focusing only on the type behavior, we ask how quickly defenders are responding to DNS abuse. That's where time-based metrics come in.

Next slide, please. Why does it matter? Because we move from actually looking at attacker to looking at defender basically for how long a domain name stays online. Next slide, please. There are different ways. There is active measurements and passive measurements. Next slide, please. The results, what this shows is basically business domains are the ones that stay the least amount of time online, followed by spam and C2 domains are the ones that stay longer.

Next slide, please. The key takeaways basically of this project is like measuring time-based metrics actually matters because better uptime metrics can hold operators accountable and make the internet safer. Next slide, please. With this, I pass to Samaneh.

SAMANEH TAJALIZADEHKHOOB: Thank you, Carlos. This was probably one of the fastest presentations we had. Thank you, the team, for all of the nice presentations you gave. The reason why it's back to me is because I know that there was a specific request from ALAC to present on INFERMAL, so we also wanted to leave time there. Excuse to everyone who has questions. Hopefully, we have some time at the

end or otherwise, we respond offline. My team will also respond in the chat now.

INFERMAL. INFERMAL stands for Inferential Analysis of Maliciously Registered Domains. INFERMAL is a project that we funded. So basically, OCTO, ICANN org, OCTO funded the project. It's conducted by KOR Labs, led by Maciej Korczyriski. We collaborated with them. The actual work is now published in a really high-length security conference. It's going to be a proceeding of that conference. And here, you see the order of the authors of the actual scientific article.

Next slide please. I guess I still need to go quite fast, so I will just skip a lot of the stuff. I think I don't need to talk about the motivation because we all know that there's a lot of DNS abuse. This specific piece of work focuses on attackers' preferences and what factors influence what the attackers choose in terms of registrars and TLDs.

Next slide, please. So, basically, the idea of this research is they wanted to disentangle the preferences of attackers. For that, the three set of feature sets are collected. The attributes were registration attributes, so what kind of attributes we've seen over the processes of registration. The proactive verifications that happen in the course of registrations and the reactive security practices. Remember, in the beginning of this session, we talked about proactive and reactive measures. In order to achieve or do a precise analysis on the relation between the factors and features

collected and preferences of attackers, there is a regression model used. I won't bore you with the details of that.

Yes, next slide, please. There are various data sets used for this study, from zone files to reputation block lists to transparency logs and many more, but this specific study is focused on phishing only. Next slide, please. So, here, in this slide onwards, I show you the set of selected features for registration attributes. There is features like free API, bulk search, what payment methods were available, pricing, discounts, whether they offer free services or not.

Next slide, please. And for proactive measures, the operational validations of contacts, the warnings and restrictions during registration processes and all kinds of other things. Next slide, please. And for reactive measures, it looked at the amount of time that the domain is up from the time that it's listed on a block list, basically, and then checking it periodically. Next slide, please. Some descriptives of the features selected. So, basically, the prices range from $0.78 to $69, nearly 50% costing two dollars or less.

Next slide, please. These are the examples of expensive domains, usps.bar or dhlcenter.net. As you can see, they kind of imitate the brand. Next slide, please. We also saw that there are discounts during the course of registration. So, basically, when you start from a specific price and then you end up with a different price because there is a discount for various reasons given. Then this is also this distribution of different free services that was add-on to registration like DNS, email, web host, etc.

Next slide, please. Coming to the results of the study. So, basically, the results suggest that if we look at the economic incentives part, we see that the price and the discounts mattered. To our surprise, price plays a less important point than the discounts that they get during the course of the process. This is for the malicious registrations. This study focused only on malicious registrations, not on compromised domains. When we look at benign registrations, we see that the mean price of the benign domains is higher, is 8.6, which is higher than malicious ones, which is 4.7. So, basically, price, whether the initial price or the eventual price, both of them are more important for attackers than for benign users.

Next slide, please. When we look at the API which allows for registering domains in bulk or to create accounts in bulk or to search in bulk, we see a strong positive relation there. API used for account creation and management. Next slide, please. When we look at the free bundles offered, even though they have a positive relation, significant positive relation, we see that they are attractive not only to malicious registrations, but also to legitimate users, which suggests that at the end of the day, they do not matter that much.

Next slide, please. When it comes to proactive upfront checks, we see that there is a negative relation. So, basically, the more checks, the less abuse. And there is a big difference between those of malicious selections and those of benign users. So, basically, this

matters a lot when it comes to malicious registers. This is a point that mitigation can make an impact.

Next slide, please. Having presented the results very fast, this is the most important of this presentation, and we actually recently published a blog on this specific part of this study. There is a lot of consideration and context to this study, which we think it's important for our community to consider and pay attention to. The blog also specifically talks about clarifications and interpretations of the INFERMAL study results.

First and most important, there are a set of factors looked into this study and are selected. It's important to pay attention that these factors are valid within the model that they are looked at, with all the other factors. What do I mean? Do you remember when I said API for account creation and management has a very positive, significant impact on the amount of abuse, 400%? What this actually means is that 400% in fact doesn't exist in isolation. It actually reflects the impact of offering the API while holding all the other factors of the model constant.

So, if you added a variable or removed a variable from the model we looked at, this would be a totally different number. It means that 400% is actually a number with the things that we measured, in relation to the things that we measured. There are many things that are not measured in this study, and in practice not everything is constant. There is a lot of things that are dynamic. So, basically what can be concluded here is that this is an important factor, the

importance, the amount of importance cannot really be taken as in isolation.

Next slide, please. Several features that are used in this study are aggregated and combined. For example, registration restrictions or consolidation of payments are an aggregation of many, many variables such as cryptocurrency, bank transfer, digital ballots, etc. So, this is something also taken into account when we are looking at the results.

Last but not least, probably the most important thing into account is that the economic implications of the results of this study is different from the actual results. What I mean is that we looked at attackers' preferences here. We didn't look at basically what they translate directly into policies or economic actions. There is a gap between what attackers prefer and what actually happens in practice. We would like to invite the community to take that gap into account when talking about the results or when trying to translate the results into policies.

With that, I finish the presentation, finish today's work. I thank all my team and everybody who worked on the INFERMAL project, the KOR Labs and all the researchers. If you have questions, you can feel free to contact Maciej, the lead of the project, or me. We also have an email which is for our team, ssr-research@icann.org, which I pasted in the chat. Thank you, Jonathan.

| | |
|---|---|
| JONATHAN ZUCK | Thank you. Let me see if I've got-- Are we over time? Well, thank you very much. I guess we'll gather up questions and get them to you offline, but we really appreciate your time meeting us on our schedule here in Prague. And with that, we thank you and thanks, everyone, for coming. Enjoy your coffee break. |
| SAMANEH TAJALIZADEHKHOOB | Thank you. |

**[END OF TRANSCRIPTION]**