ICANN71 | Virtual Policy Forum - Plenary Session: Understanding Reputation Block Lists
Thursday, June 17, 2021 - 10:30 to 12:00 CEST

BRENDA BREWER: This session will now begin. Please start the recording.

[Recording in progress]

BRENDA BREWER: Hello. And welcome to ICANN71 plenary session, understanding Reputation Block Lists.

My name is Brenda Brewer, and I am the remote participation manager for this session. Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior.

During this session, questions or comments will only be read aloud if submitted within the Q&A pod. I will read them aloud during the time set by the chair or moderator of this session. Interpretation for this session will include English, Chinese, French, Russian, Spanish, and Arabic. Click on the interpretation icon in Zoom and select the language you will listen to during this session. If you wish to speak, please raise your hand in the Zoom room. And once the session facilitator calls upon your name, our technical support team will allow you to unmute your microphone.

Before speaking, ensure you have selected the language you will speak from the interpretation menu. Please state your name for the record and the language you will speak, if speaking a language other than English.

When speaking, be sure to mute all other devices and notifications. Please speak clearly and at a reasonable pace to allow for accurate interpretation.

All participants in this session may make comments in the chat. Please use the drop-down menu in the chat pod and select respond to all panelists and attendees. This will allow everyone to view your comment. Please note that private chats are only possible among panelists in the Zoom Webinar format. Any message sent by a panelist or a standard attendee to other standard attendees will also be seen by the session's host, co-hosts, and other panelists.

To view real time transcription, click on the closed caption button in the Zoom toolbar. And with that, I will hand the floor over to LG Forsberg. Thank you.

LG FORSBERG: Thank you, Brenda. My name is LG Forsberg, and I am your moderator for this ICANN71 plenary session titled: Understanding Reputation Block Lists.

For those of you who don't know me, I'm a Swedish domain name industry veteran with experience from both the registrar side as a technical product manager and registry liaison as well as from the registry side in various technical and channel-related roles.

I am currently the CTO at iQ, a commercial provider of services for the domain name industry like the iQ Abuse Manager and domain analytics.

When I'm not integrating the next Reputation Block List into Abuse Manager, I also consult for various registries and registrars in technical and policy matters.

Last, but not least, I am the founder and curator of Nordic Domain Days, a yearly Scandinavian domain industry conference, who really can't wait for face-to-face meetings to be a thing again.

During the session today, we will engage the crowd. And we have prepared three poll questions for you to answer while listening in on the discussion. Two of them will be presented to you now.

While the last one will pop up just in time for the last segment of today's agenda.

If you would please present the poll questions.

BRENDA BREWER: Thank you. On your screen you should see the first polling question. Are you familiar with the following types of security threats? Please check all options that apply.

The options include spam, phishing, malware, pharming, botnets, or others. Once again, are you familiar with the following types of security threats and check all options that apply: Spam, phishing, malware, pharming, botnets, others. And we will close the poll in about five seconds.

Thank you.

We can close the poll at this time, please.

And may we have polling question number 2?

Polling question number 2: Has a domain name you managed ever experienced one of these security threats? And you may answer yes, no, or I do not know or I do not manage domains.

Sorry, there's four options. Please choose one. Again, has a domain name you manage ever experienced one of these security threats? Yes, no, I do not know, or I do not manage domains.

And we'll give it about five more seconds, please.

And thank you very much.

We can close the poll.

LG, back to you, please.

LG FORSBERG: Thank you, Brenda.

I will continue by sharing a brief introduction to the subject before we continue on to the panel section. Today we have a number of parties listening in that needs to associate with Reputation Block Lists in different manners. These parties, I would say is ICANN, who provides statistics within DAAR, using Reputation Block Lists, and they also require new gTLDs to at least monitor their (indiscernible).

We have contracted parties in the form of registries and registrars who might be using one or more Reputation Block Lists to monitor their TLDs, or their list of domain names registered in case of registrars.

We might have service providers like hosting companies, and they might use Reputation Block Lists to keep their customer emails free of spam or to make sure that websites that they host are not being reported for bad things, such as phishing or malware.

Lastly, we have end users which for this discussion are defined as someone that is the registrant or the operator of a domain name and/or a website which is, I guess, the one group here today that might not want to have a relationship with a Reputation Block List.

That said, Reputation Block Lists will, of course, mean different things, depending on who you identify with in these groups. For ICANN, a registry or a registrar, it will most likely be a tool you use or a partner you work with, while for an end user it might spell disaster if you end up on one with your email sent to the black hole of the junk folder or you get the red screen of death from Google showing up in Chrome instead of your website.

Also, as the topic might have told you, while Reputation Block Lists is integral for the DNS abuse discussion, we're not here today to have a 90-minute discussion on what the definition of DNS abuse is. Instead, we are here to gain a little bit more understanding about how Reputation Block Lists do their work and how we, the audience, the contracted parties, the service providers, and the end users can work with them.

Now we'll be hearing a lot about Reputation Block Lists today, but a simplified brushup in front of that discussion can't really hurt.

So a Reputation Block List is a collection of indications or reports offered to help stave off abusive behavior in one way or another. The reputation part of the name often signifies that there are some shades of gray to their way of determining if a domain name should be listed or not. It's often not a simple yes or no.

An example of this is the Spamhaus implementation of reputation, with each domain name starting out at zero, gaining negative points for doing good things and positive points for doing bad things.

Once enough points have been accrued -- for example, five or ten points -- the domain name is listed in the blocklist.

The block word in the name pays a little bit of homage to the origin of these providers where they were being implemented or are being implemented in email servers and firewalls to stave off spam or unwanted traffic. The lists were and are to block things.

So mitigation of abuse -- sorry. Of course.

So mitigation of abuse might not be the only use case or sometimes not even the primary use case.

Some Reputation Block Lists will offer up a full list of information on how a certain bad resource was discovered, when others just give you a hint that something bad has happened on a particular domain name.

Something that most reputation lists seem to have in common is that understanding what they're doing, how they're doing it and what their data is meant to help with is not always easy, especially if you're an end user without any prior connection to these types of resources.

But to try and get to the bottom of how Reputation Block Lists gather their data, categorize their content and really dig into how this machine crunching, in a way, occasionally spitting out a new list, go about their things, we have gathered a few representatives from the Reputation Block List here today, and we're also joined by a group of representatives from the parties presented earlier to talk about their experiences with Reputation Block Lists.

And with that, I would like to introduce our first participant today and let him introduce himself and the company that he works for.

Please, Mr. Carel Bitter from Spamhaus.

CAREL BITTER: Good day, everyone. My name is Carel Bitter. I work for Spamhaus, and I'm in this capacity mostly involved with the domain reputation side of what we do. We make multiple datasets, but today let's focus on the domain part of what we provide.

I've been involved with domain reputation for over ten years now. And, yeah, I've been a previous participant in a whole bunch of ICANN sessions and whatnot.

So happy to help and any -- answer any questions that people might have.

LG FORSBERG: Thank you, Carel.

Please, Mr. Roman Huessy from abuse.ch, if you would please introduce yourself.

ROMAN HUESSY: Yes, I'm Roman Huessy, and I'm the founder of abuse.ch. Abuse.ch is a nonprofit project which is run at the Bern University of Applied Sciences, and the goal is to get their information about botnets, and other, and publish their corresponding information for free for everyone.

LG FORSBERG: Thank you, Roman.

And our last participant from the Reputation Block List today, Mr. Ben Coon from WMC Global.

Please introduce yourself.

BEN COON: Thank you. Ben Coon, WMC Global. We are a -- primarily a -- run a phishing platform. We provide blocklists on phishing URLs to SMS providers and firewalls and people who are getting hit mainly with credential phishing over a multitude of lures.

Happy to answer any questions than anybody have. Thank you.

LG FORSBERG: Thank you, Ben.

To continue with our discussions from the parties, I would like to start with introducing Samaneh Tajalizadehkhoob from ICANN. Please introduce yourself. And if you have any leading points that you would like to bring into this discussion, then now is the time.

SAMANEH TAJALIZADEHKHOOB: Hi, everybody. I'm Samaneh Tajalizadehkhoob. I work at ICANN's office of CTO in the security, stability, and resiliency group. I'm also the project lead for the DAAR project and have previous experience in academia, working with reputation feeds.

So today I will be talking more about how, as an ICANN organization, we view and use these feeds and in what projects.

Thank you.

LG FORSBERG: Thank you, Samaneh.

And then I would like to introduce Mr. Matt Thomas from Verisign. Please introduce yourself and bring any leading points that you have to the table.

MATT THOMAS: Thank you. My name is Matt Thomas. I work at Verisign in the Cybersecurity Strategy and Research Department as a distinguished engineer.

I'm also an ICANN SSAC member and also currently serving as the vice-chair for the board of directors at M3AAWG, and I'm looking forward to the panel today to discuss the various use cases of RBLs.

Thank you.

LG FORSBERG: Thank you, Matthew.

From the registrar side of things, we are joined here today by Reg Levy from Tucows. Please introduce yourself and bring any leading points you have.

REG LEVY: Hi, everyone. My name is Reg Levy, and I work at Tucows as the head of compliance. I also sit on the Registrar Stakeholder Group's DNS Abuse Group. Actually, I'm co-chair of the DNS Abuse Group.

Good morning. I'm still waking up, apparently.

I'm looking forward to this discussion.

LG FORSBERG: Thank you, Reg.

And last but not least, we have, from the At-Large community, Joanna Kulesza. Please introduce yourself and bring any leading points you have to the table.

JOANNA KULESZA: Thank you, LG. Thank you for having us. Thank you for having me and giving me the opportunity to speak on behalf of the end users.

I am the co-vice chair of the At-Large Advisory Committee. I focus on capacity-building, and I do believe there is a capacity-building component to this discussion, this panel, and Reputation Block Lists.

As a leading point, just a disclaimer, by the poll itself and participation from the At-Large community, it is not true that the end users are not interested in Reputation Block Lists. Quite to the contrary. And I will try to get that point across as we progress in the panel.

Excited to be here. Thank you for having me.

LG FORSBERG: Thank you, Joanna.

We will then step from the introduction section of this session into the first interview section where we talk to the Reputation Block Lists and try to increase our knowledge of how they do things.

I would like to start this section off by asking Carel how you yourself define a Reputation Block List.

CAREL BITTER: That's a good question. Let's start with the blocking part.

I think in many cases, it is not just useful blocking. And if you look at the whole registry/registrar perspective where registries or registrars are looking at the data to find problematic domains or problematic customers, then there's not -- they're not really involved in any blocking per se. It's more of a remediation-type use case.

So personally, I usually talk about our data sets and not about our RBLs, especially in a context like this. So -- and I think that's a really important thing that you need to take into account whenever you start working with any data set. What is it meant for? And what are you doing with it? And are you using it as intended? Are you using it in the way that it was -- what it was designed for?

And if not, I mean, you should be -- I'm not saying you cannot use something that's meant to do one thing for another thing. But you need to be aware that your use case may be a little bit different to go on to one part.

As you mentioned earlier, about how our systems work where we actually assign scores to domains, the higher a score -- a domain reaches in terms of scoring for us, the more comfort we get that something really -- something bad is going on. So this allows people to act differently on different gradations of scoring.

So I think, yeah, you really need to understand the data that you're working with. And if you don't, then you need to reach out to the people who create the data and say, Look, I don't really understand this part or I'm trying to do this thing with it, is that a smart idea? Certain things you might want to block in the context of an email but probably not in the context of the DNS level. A good example would be shorteners. Shorteners in email are a very well-known sign of a problematic email. Any legitimate email usually doesn't include shorteners.

But to block shorteners at a DNS level is probably a little bit of a harsh treatment for your end users. So there's definitely lots of things to consider there.

LG FORSBERG: Thank you, Carel.

I would like to follow up with two follow-up questions specific to Spamhaus. These are: What are the typical types of evidence that Spamhaus would consider strong enough to list a domain name in a Reputation Block List? And is this evidence in any way shared with the adopters of your data sets?

CAREL BITTER: In our case, it's not always easy to share the evidence because all of the data we get is -- comes from ISPs and other networks where they have an agreement to share data with us, but we cannot always share things on. In case where people doubt things, we're always open to questions.

In terms of the registry/registrar use case that we work with, there is a support channel for people to say, hey, we see this as listed and we don't really understand what's going on here. Is it maybe a false-positive? Is there something else going on? We can always look at it, and it may be that we have something we can share. But it's not -- it's not, like, default shared with. Like, if you get a list of, let's say, 100 bad domains, it doesn't contain also 100 spam samples, for example, or 100 binaries of malware. That's basically a case-by-case basis where we need to determine what we can share and whatnot.

LG FORSBERG: Thank you, Carel.

Roman, do you -- what would you say is the difference in this subject with abuse.ch?

ROMAN HUESSY: Sorry. Can you please repeat the initial question? I was just answering the Q&A pod.

LG FORSBERG: I'm sorry. No problem. So what is the typical kind of evidence of abuse that abuse.ch would consider strong enough to list the domain name in your data set?

ROMAN HUESSY: That's a good question. Thank you.

So usually the type of evidence is malware that is being served on compromised websites, so domain names or domain names that has been registered by a threat actor.

And what that means is that the system checks actually every submission whether there is a malicious payload being served. And if this is the case, the evidence is actually being published on the project websites so everybody can use it. And then an abuse report is being sent to the corresponding hosting company.

So the project is pretty transparent in terms of providing evidence, as every evidence is being shown publicly on the website and everybody can actually check why something is listed.

LG FORSBERG: Thank you, Roman.

For the understanding of a data set, would you be able to describe the life cycle of a generic report, that's in your case malware then, from first -- the first time of it being submitted to you until it goes away?

ROMAN HUESSY: Sure. So speaking of that, malware, it's, more or less, an easy topic unlike phishing domains or unlike spam domains as I expect a certain kind of response from the remote host. So that means I can easily check in an automated way whether a website is still serving malicious content, so to say. And I do that in an automated way. And once that malicious content disappears, the website or domain name is being flagged as offline automatically. And that then actually means the domain name or URL automatically disappears from the blocklist that I provide.

And because of that, the good thing with that is that URL only stays listed as long as it actually serves a threat. And whether it actually still serves a threat is being checked several times per hour, usually every ten minutes.

So once an end user fixes the threat, the domain or the URL will automatically disappear from the blocklist within an hour usually.

On the other side, it actually means that if the problem was not really fixed, rather than, for example, the malicious content has just been deleted but the root cause has not been fixed, means the threat actor uploads the malicious content again, the system automatically check whether the malicious content is there again. And if it's there again, the domain name or URL will automatically be listed again.

LG FORSBERG: Thank you, Roman.

Ben, would you say that Carel and Roman has described here differs significantly from what you at WMC Global see in terms of phishing reports, their life cycle, and the evidence of phishing?

BEN COON: I would say that for WMC Global, we follow mainly the same model that Roman was talking about.

We verify maliciousness on sites before they get listed. When we provide a blacklist to, say, like, an SMS provider to stop smishing, they usually only use that list for the last 24 hours because the phishing sites will go offline. We will relist a phishing site if it comes back up. We have automated checks that go out, and they will test the life cycle of how long that phishing URL is live.

The other thing that differs in what we offer and what some of the other blocklists offer is we don't list domains; we list out the full URL. That way you can see exactly where the malicious content is. And then, again, once that malicious content goes offline, it will be delisted unless it then comes back online.

LG FORSBERG: Thank you, Ben.

Back to Carel. You mentioned earlier that a user of a Reputation Block List needs to be aware of what the blocklist is meant to be used for.

What would you say is the primary use case for Spamhaus? And so which is the primary user you are trying to serve? And would you say that this has changed over, say, the last ten years?

CAREL BITTER: Obviously since we have "spam" in our company name, we are coming out of the email corner of things.

The data set is -- so it's a little bit more complicated. The data set that we publish publicly, the domain blocklist, is our view of what domains are currently, have a bad reputation. It's mostly used for email, but there are also malware domains in there. There are also botnet metacontrols in there. There are also phishing domains in there, all the domains that we think are bad on one big pile.

The distinction is done through, in that case, DNS return codes. So it's possible to segment out only one part, like only the phishing or only the malware.

So -- but what we do is we provide data. We have little control over how people use it. So the format we publish, the public free version, that's free for everyone to use, is a format that is mostly used within email servers. So it's a DNS BL. You do a DNS query, you get an answer, and you work with that answer to do something within whatever your local policies are.

There are some versions available that work for specific threats and specific use cases. So there are different subsets available, for example, to be used on DNS resolver level. There are versions available to be used for registries and registrars to do investigation and remediation.

So there is -- what's the primary use case? I guess, the primary use case is people who want to make reputation decisions based on a domain name.

I wouldn't say it's for email only. That may have been in the past, but it's certainly not where things are now.

Most domain names, as Roman will be able to tell you, that are used for malware and botnets, you will never encounter in an email. So you can check all your emails for these domains, but you will just never see them. They are being used at a different level where an infected computer will reach out to a certain domain name to contact its manual control center to download additional malware.

So the place where you would check for those domain names is not at your email server, but it would be at the DNS resolver or in an IDS or anything like that.

So the use case really depends on what kind of problem you're trying to fix.

As what Ben said, for SMS checking, for example, again, a lot of the domains we see used in SMS or shorteners that redirect in some domain, you never see those in email.

So I think it's a lot wider than just email. If you have a security issue and domain names are involved, then there are data sets like what we produce, like what Roman produces, like what Ben produces that can be helpful to solve your security issues or to address and become aware of things that go on.

LG FORSBERG: Thank you, Carel.

I now have a question for all three of our Reputation Block List providers. We have heard all of you talk a little bit about how you receive and collect reports or indications. And you can tell from this that we are talking about mostly reports or machine-based detection crawling the Internet and so on and so forth or looking at emails that you have access to.

Would you say that there is ever a form of human investigation in the -- in your company, in the Reputation Block List? Let's start with Ben.

BEN COON: Yes. We use -- we use a host of threat hunters that will double-check, or spot check a lot of the credential phishing that we see. We will also take a look at anything that does not score high enough to automatically get listed into our Reputation Block List before it gets put in.

And then any type of false-positive that we encounter will go back to the team. The team will do an investigation and either manually add it or manually remove it.

But I would say that we have quite a bit of human intervention.

LG FORSBERG: Thank you, Ben.

Roman.

ROMAN HUESSY: Yes. So speaking about URLhaus, in particular, a project that tracks malware sites, it's a community-driven project. That means that the data sets that I produce is just one part of the data set. The other part is generated by the community.

And when we speak about community, it's important to say that there are two types of reporters. One type are the stranger reporters, so to say. So these are users that I don't know that I don't trust. Whenever they report something to the project, it is manually being vetted.

On the other side, we have trusted reporters. And trusted reporters means that if a trusted reporter reports something to the project, the site automatically will be listed. However, the system, of course, will still check whether the website is serving any malicious content or not. But the URL will go directly into the database.

However, if you now use it in a blocklist approach, which is just one use case for the data set, of course, the site will -- or the domain name will only appear in a blocklist when there is actually malicious content that is being served.

So to answer your question in short, it is a mix of manually vetted URLs or domain names and URLs that are vetted by a machine.

LG FORSBERG: Thank you, Roman.

Carel?

CAREL BITTER: Yeah. In our case, it's a mix as well of automation and human investigation. If you want to assign a reputation to every domain name that exists on the planet, then obviously you will need automation. There is just too much -- there are too much domain names around, and there are too many new domain names every day to have humans look at everything and investigate everything.

So part of it is automation, but there's definitely a human component to it as well where our investigators look at things that are suspicious or things that are scoring just close enough to not get it or things where you say, okay, I would have expected this to score higher and it doesn't. So it's always the combination of human and machine intelligence going on.

LG FORSBERG: Thank you, Roman.

We are running out of time for this segment of the session. However, there's a popular question that is the next one that I would bring to the panel. How likely would you say that there is a report or an indication that is a false-positive? And what is the most common reason for a false-positive? Let's start with Ben.

BEN COON: I would say on -- within data that we see that is being reported on blocklists, there are always some false-positives. It's our drive to get those false-positives as low as possible.

I would say that mainly the main reason I see for false-positives are lack of understanding of what is malicious and what is not malicious or people who will get something that they just don't like, and they will report it as malicious, even though it is not.

LG FORSBERG: Thank you, Ben.

Roman.

ROMAN HUESSY: Yeah, so there are, of course, false positives. There always have been. And they can be reported through the web UI by everyone.

If a trusted reporter, for example, starts to feed data that is known malware -- for example, phishing or spam sites or other stuff -- he will be simply blocked and excluded from the project.

So the stream of data is, in my opinion, pretty good. But as I said, false positives, they may happen.

LG FORSBERG: And lastly Roman. No, lastly Carel. Excuse me.

CAREL BITTER: Yeah, false positives will always happen. And like Ben said, it's our job to make sure the number is as low as we can make it.

As to what people report, there is -- you know, many years ago when email providers started having "this is spam" buttons in their whitemail interfaces, one of the biggest things that people reported that actually turned out to be not spam or not malicious were mails from -- mails that they just didn't like, like an invoice that they didn't want to pay or a rejection letter from job application. So in that case, most of the systems work in such a way that if enough people, if a certain critical mass reports something, then automation kicks in, in the signs of bad reputation to an IP or domain name or a sender or a beacon selector or whatever. And there's always the chance that especially if user reports are a part of the stuff that you do that people will report things that, like Ben said, that they just don't like, then it may not be spam and it may not be malicious and it may not be phishing. But for a lot of people, the distinction between the delete button and "this is spam" button is nonexistent.

So, yeah, that's always a concern. But, yeah, obviously if you think about it, and I guess this goes for Roman and Ben as well, is the fact that any false positive that we will have in our datasets, and there will always be issues. You may have a week where nothing bad happens, and then the next week you have like three or four or five or whatever. But it's a fact of life they will happen. The important thing is how you deal with it, get them out there quickly, you know, make sure that things are being dealt with properly.

But, obviously, as providers of the data, as the creators of the data, we have an interest ourselves to make sure that our data is as good as it can be. You know, if our data is not good, the people will just stop using it, and then at that point, what's the point?

For the creators of the data, and I think especially, you know, Roman will probably agree with this, the data gets more powerful the more people start using it. If there is a domain name that is being used for -- for like a botnet command-and-control type scenario, the faster that domain name gets, you know, taken out or the more people that block it, the safer the Internet gets.

And if we -- or Roman reports it and people say these reports are no good, then that sort of is opposite of the mission that we want to have.

What we want to have is the data that we produce, we want to have it used as widely as possible and we want people to take reports as good as they can and so that in the end it becomes a safer place. Because that's, I think for all three of us, for Ben and Roman as well, that that's what drives the work that we do. We are trying to solve security problems for networks, for registries, registrars, for end users, in the end. In the end, it's all about protecting the end user. And the more people that use the data, the more end users will be protected. And people will not use the data if the data is not good.

So, I mean, this is a very important thing that we always want to do.

There's infinite ways to be more aggressive and where you say, okay, if I do this, I will block more -- I will catch more bad things happening. But there are many scenarios where there's always someone somewhere who is doing like almost the exact thing as the bad guy or has the same sort of weird URLs or has the same sort of strange host names. Like I see this host name; it must be bad. Yeah, okay, 99 out of a hundred are bad but the last one is actually not.

So it's always a balancing act, and you just need to be sure that you try to be as good as you can be and have a good process of dealing with any issues that arise.

And I've already typed it in one of the Q&As, but, for example, with us, anyone can go in and remove a domain name. You don't need to -- you don't need to fill in lengthy forms and you don't need to contact us by phone or whatever. It's a self-serve thing over the Internet like Roman's button.

It's obviously in our advantage to make sure the process of dealing with any false positives that might occur is as smooth as possible for whoever is the victim of a false positive.

LG FORSBERG: Thank you, Roman. Let me interrupt you there.

We are in the audience question segment, and I would like to do one more before we move on to the next segment of the session.

I would like to ask Roman how -- if Reputation Block Lists or threat intelligence feeds that have -- let's call it original data, ever cooperates with one another? So, for example, would -- would abuse.ch provide, say, Spamhaus with data on false positives or important notices that you might have?

ROMAN HUESSY: Yeah, so the question has raised in the Q&A pod before, and I think it's a very important question.

So there are, of course, processes where information about current threats are being shared with vendors or RBLs. Some of these processes, they are bilateral. They exist, for example, with Spamhaus or with Safe Browsing or with other RBL providers. But as -- in terms of abuse.ch, the dataset is openly available for everyone. Everyone can consume it. And there is a large portion of commercial threat intel providers, for example, that are consuming these feeds and doing with that whatever they want.

And as there is no registration needed, there is, of course, the issue on my side that I don't know who is using my feed as it's public domain.

On the other side, regarding the quick question of exchanging information about false positives, I think this is an important topic. And as far as I am concerned or aware, there are no mechanisms in place at the moment how to report a false positive.

How I handle that is when something gets flagged as a false positive, it of course -- it of course gets out of the feeds. And I would expect that other RBL providers or other -- other threat intel providers, they will notice and remove the entry as well. But of course I don't have any influence on that. And I think that's something that would be -- it would be a topic that needs to be discussed, to have an exchange platform where false positive -- information on false positives, for example, can be -- can be announced and shared with other providers.

LG FORSBERG: Thank you, Roman.

We have now reached the ICANN presentation of this session. So I leave the word to Samaneh.

SAMANEH TAJALIZADEHKHOOB: Hi, everybody. For the record, I am Samaneh Tajalizadehkhoob. I represent ICANN's office of CTO today, and I will talk about Reputation Block Lists and our view on it.

Thank you. Next slide, please.

So, so far, we had very good discussions. I think some of the points are already made in the presentation is already covered by the panelists and also by the questions that are answered. Excuse me if there is repetition.

I wanted to start from the base what is a Reputation Block List. They can come as an IP blocklist or domain-based blocklists. What they generally represent is entities that are regarded as malicious, untrustworthy or simply bad reputed. I listed some of the use cases that are there in the industry or research or academia. And people use it to feed DNS firewalls, for preventing malicious traffics. It's also used for filtering out unwanted traffic, which is typically spam and phishing emails. It is used by CDNs to prevent delivery of malicious content to their customers and also used as part of incident response or law enforcement purposes for identifying malicious infrastructure involved in the attack.

They come in different sharing mechanisms. So I think today we have some examples of the ones that are rather open source, but they can also come as commercial, either available through a rate limit, license based, or pay-per-use, and can be -- are normally maintained by for-profit companies specialized in threat intel.

There are open source ones out there that are mostly used by academics, and also others. Examples are Spamhaus, abuse.ch, Phish Tank, the one I am familiar with, among several others.

They can also be threat specific. So certain lists are only focused on certain threats. For instance, abuse.ch, maintained by Roman, have different feeds which are focused either on botnet or on malware, on ransomware, I think that one is stopped now, et cetera, or they are general, more general. They contain all kinds of threats, which an example is SURBL.

Next slide, please.

We've already talked and discuss this extensively in the panel. The title slide says characteristics and drawbacks of the list. I want to focus on both aspects because something that for a researcher is a characteristic can be a drawback for another researcher.

So the main point of using an RBL from our perspective is it's important to understand what you want to do. And actually more important to understand what is the methodology behind an RBL and what does it represent and how you can use it best to your purr.

Some of the characteristics are that some lists can be overspecialized, so -- because they are geared towards a specific purpose. Any person who wants to use them has to understand that purpose and see if it fits their purpose.

Some -- Typically -- and I am making these points generally speaking, as a researcher. Typically they have limited coverage due to the physical limitations and vantage points. Normally the data feed providers are located in certain geo locations, and so they may have a lower representation on certain geographical location.

I know that over time, since I have experienced working on this for already seven, eight years, over time most of the lists have been improved in terms of coverage, but still it's important to have this -- to have this in mind when doing any kind of analysis on the lists.

A small point is that the idea of looking into quality and different reliability metrics on the blocklists is not new. It has been already researched already since 2009, 2010. So in this presentation, I'm listing some of the academic and industry research that has been done on the lists, but you can contact me for more if you need it.

There is another issue or characteristic of the RBLs is that there is typically limited documentation on the internal methods. One reason that I understand from talking, having personal conversations with the list providers is that it is not straightforward to document the whole process since some of it could be ad hoc or it could be very detailed or reactive. So it's hard to maintain a real-time documentation that represents everything that is -- that an RBL provider does to a specific feed.

There is also -- Since we have so many RBL providers, expectedly there are different -- there are a lot of varieties in terms of methodology, so -- in terms of data collection, curation, maintaining and labeling of the blocklists. This leads to different effects on coverage, reliability, effectiveness, and speed of reporting.

Important is that this is not necessarily a bad thing. This could actually be a good thing. The diversity could bring more information into the table as long as the user understands the diversity and addresses it correctly.

Next slide, please.

I've made this point already. So what -- why is it important to know these drawbacks, aka characteristics? First of all, it's important for every user, let it be network operators, researchers, security companies, that are relying on these resources to know the differences between the lists and the drawbacks, the caveats, and also knowing them to design more effective defenses and curation methods as well as having had these in mind while doing research or reporting results.

A good example of this that I personally come across a lot in the ICANN community is that we often say there is this research that showed phishing is going -- the trend of phishing activity is going up. Why does it contrast the other research that person B is doing that shows the trend is going down?

So this is the exact point that I'm trying to make, and I think this is the main point of my slide set, is that depending on what dataset that is being used, depending on what time frame that is being taken into account, what labeling process that the dataset acquired, et cetera, a trend can be changed.

So there is no absolute trend in any of the feeds, at least from my experience, that one can report this is the absolute trend. It's always partial, so it's always from the perspective of the feed. As Roman and Carel and Ben already reported several times, that from our view of the RBL, this is what we see. So it's always partial, and it always depend on the methodology.

Next slide, please.

In the next few slides I will highlight some of the use cases that we have within ICANN. So how we use RBLs.

DAAR is -- the Domain Abuse Activity Reporting is a project -- is one of the main projects that use RBLs that most of you may be familiar with.

I won't go through the details of what DAAR does, but in summary, we -- so the system takes domain names from registry zone files and then takes domain names from the set of preselected feeds for phishing, malware, botnet command and control and spam as a delivery vector. Then it overlaps the domains from the zone files and their RBLs and processes, calculates, and creates metrics, daily and monthly metrics for different kind of analysis that you might be familiar with from the ICANN monthly -- DAAR monthly reports, which are public on the ICANN websites. They show trends of where DNS security threats are concentrated per point in time and how this concentration is changing over time.

So what to have in mind, that the system, of course, already does extensive pre-processing and cleaning and unifying the RBL data feeds that it employed given all the corner cases, et cetera. The details are listed in the methodology document of the project.

Next slide, please.

BRENDA BREWER: Excuse me. Before you continue, could I ask that you kindly slow your speech for the interpreters. Thank you.

SAMANEH TAJALIZADEHKHOOB: Sure.

We have also done a project. As the SSR research group we have also done one project for the ICANN Compliance Support in which we created a snapshot metrics for registrars. So these are not -- this is not a -- this was one-time project so far in which we did a similar thing to DAAR but only for registrars. This will focus on phishing and malware for a specific period of time, so it was just certain amounts.

And we calculated metrics for registrars and registrar families showing threat concentrations in one point and over time.

For this specific project, we had access to BRDA since we can only use that for compliance purposes for now.

Next slide, please.

Other research projects that are ongoing at the moment within OCTO, we use RBLs for developing predictive models, to be able to predict when a domain will become malicious and also to extract patterns to characterize malicious domains.

We are also planning to employ a similar method as it's been used by the COMAR study that has been presented in the previous ICANN meeting Tech Day to distinguish between maliciously registered and compromised domains using the RBLs as one of the inputs.

Next slide, please.

So how currently we -- or for most of the projects that I listed in the previous slides, we evaluated -- or we did evaluation on the selection of the reputation list, as well as monitoring feeds for a period of time before including them as part of our research work.

What we did was to basically pick the most reputed lists within academia and industry based on publication. So reputed based on publications because it can be subjective. And we opted for those that have better documentation and data standardization, recording of our processes and complement to our existing standards in terms of coverage.

Yeah.

Next slide, please.

However, we are planning to move to a more robust evaluation criteria. This is an ongoing work.

And, again, I would like to emphasize that doing work on RBL evaluation is not new, has been done. What we are doing here is just to make it more complete and to make it relevant for the feeds as they are today.

So we are working on developing metrics for what we call purity, which is the false-positive/false-negative rates of the list. In parentheses, these are not straightforward metrics given that having ground-truth data, which is the data that is manually labeled within this set, is not straightforward; is actually the first issue of any researcher who works with RBLs.

We are looking into estimating the coverage of the lists, the responsiveness or time to respond, accuracy, how detailed information a feed is, how stable they are over time, and how liveliness, so how much of the listed domains are truly positives and active when they appear in a feed.

These are not fixed ones yet. We are trying to work on them to see if we can achieve reliable results. They might change over time as we're working on the research project or they might stay the same. But we will keep you posted about this work. As for now, it's just work in progress.

Next slide, please.

Here are the references that I used during my presentation. And so the presentation is over now.

One last point that I would like to make because as ICANN Office of CTO, we are doing several trial-and-error projects using RBLs and specifically for registries and registrars to try to provide more information in terms of security threats.

Some of the reasons why developing certain metrics that we've talked about previously in the community is not straightforward and we do not want to publish something that is not reliable is the exact issue that we are discussing today.

For instance, we've previously talked about uptime, as in the time it takes for operators to react on certain security abuse. If I connected it back to the discussion that was happening in the panel but also in this presentation regarding difference in methodology behind each list, then one can quickly realize that if we take several lists and try to create a metric for a certain operator that represents uptime, we are mixing different methodology. And the metric by itself doesn't represent something that is reliable per operator. It also depends on the methodology of the RBL provider.

So these are examples of certain difficulties that we have in our projects, that we are trying to address them. But it's good actually to bring them -- bring this up in the community discussions, and this was a good opportunity.

Thank you, all. I'm happy to answer questions, if there are any. I don't know about the time.

LG FORSBERG: LG Forsberg here, moderator. We are going to have to move on to the next section of this session due to time constraints.

And the next session is the concluding discussion where I would like to turn to our contracted parties amongst the panelists, so that means Reg and Matthew.

I would like to ask you both if you use Reputation Block Lists today and, if so, how. Let's go with Reg first.

REG LEVY: Thanks, LG.

At this time, Tucows does not pay any retail blocklist for their data. However, we do receive reports from many of the big names on a regular basis.

Unfortunately, most of the reports tend to not include the full URL, so it's difficult for us in every case to mitigate because we can't necessarily distinguish exactly what is the issue in the sense of is it a compromised domain or is it a bad registrant?

In some cases, we can. We can take that domain, look at it in our system and say, okay, yes, this was also part of a fraudulent credit card purchase. In those cases, often the domain has already been taken offline before we get the report. But we still get the report.

So at this time, we don't. We're currently looking at some of the options available. Important to us include speed of removal from the blocklist. So it's not enough to me to just report that a domain is bad but also for me to say, great, thanks so much for the report, we've dealt with it, and then to have them remove it from the list.

LG FORSBERG: Thank you, Reg.

Matthew?

MATT THOMAS: Yes. I'd just like to make a few additional points around the usefulness of RBLs that have already been made by some of the panelists, and that is more around the concept of context which Carel started to bring up and was also mentioned by other panelists.

I think it's very important to understand the context in which RBLs are being applied. We should start off by saying that they are very effective tools when used in the proper context. And I think a lot of that comes back to understanding the genesis of what these RBLs were originally designed for. They were designed as tools to protect end users and networks and environments.

So in that kind of a use case, they have certain properties that are intended for that purpose. They might be a little bit more liberal in terms of -- excuse me -- what they include in the RBL. But they might also be a little bit more conservative in terms of removing some of those entries.

So from the perspective or use case of enterprise security as a security practitioner, that is a good property, right? And in that scenario, you're less concerned about the false-positives and you are more concerned about protecting those end users.

But when you start to use RBLs in other contexts, it's important to understand the properties of those RBLs and how those will influence the various different ways you use those RBLs in different measurements, tools, or studies. Knowing the nuances of how those RBLs are constructed, how they're maintained, how they're operated, how they're vetted will ultimately influence your ability to use that RBL in any other kind of systemic manner.

So to that end, I think it's important that we all just stay aware of how these are used and apply them accordingly.

Thank you.

LG FORSBERG: Thank you.

REG LEVY: I want to briefly underscore that, that the transparency in understanding how the RBLs are constructed is extremely helpful to those who are using their data.

LG FORSBERG: Thank you, Reg.

Again, Reg and Matt, would you say that there is more that contracted parties, registries, registrars, could do with the help of Reputation Block Lists if any particular point of their -- the data was given to you or updated in another way?

MATT THOMAS: This is Matt Thomas for the record.

That's a really interesting question there, LG. And I think what strikes me in that question is the word "more." What does that mean exactly? And how do you measure "more"?

Does "more" in the terms of domain abuse context mean that more domains are taken down? Or does it mean that we're more effective at disrupting and dismantling the underlying infrastructure tooling that is supporting the DNS abuse types that go out there.

So I think it's almost that we should take a step back that it's not just the RBLs and the contracted parties in this ecosystem. It's a broader community that needs to focus on DNS abuse. There are numerous different entities out there that holistically would need to work together to actually make this more achievable, right?

We need to involve the hosting providers, the CDN, the mail providers, law enforcement. And holistically working together, we can start to, you know, combat the abuse and do more in that way.

I'm sure the RBL providers could give you tons of examples where they have collected data, they've identified a phishing domain, it was placed on the RBL, it comes down, and two hours or a day later it comes up on a couple new domains. So did that actually abuse remediation do effectively more? Or did it just cause name diffusion and push that domain abuse somewhere else, right? Is the actual abuse itself being disrupted?

And so the problem underlying is that we need to be going after the infrastructure and the systemic things that are supporting these.

Another great example would probably be DGAs and botnets. I mean, take a look at, for example, Conficker, right? It's a decade old and it's still going on, or Avalanche. Efforts are still having to be made to combat these types of DNS abuse because the underlying problem that the host machines aren't being patched and the underlying systems aren't being remediated is the crux of the real problem here, right?

And so I think there is an opportunity to the latter portion of your question there, LG, of what could be done with RBLs from a contracted party's perspective a little bit more is that I think the RBLs are actually positioned very well in the ecosystem based off of their unique observation space where they can get telemetry data to help inform how effective the efficacy rate is of dismantling these underlying platforms that are systemically supporting domain abuse.

So I hope that we can work together with the ICANN community to push in that direction. Thank you.

LG FORSBERG: Thank you, Matthew.

Reg, do you have additions?

REG LEVY: Yeah. And I agree with Matt's framing of, you know, what is "more" in this circumstance.

We are working constantly with various blocklist providers to establish whether or not we can use their services because of the false-positive mitigation as well as removal post-mitigation on our end because as some of the -- some of the conversation earlier indicated, there are a number of false-positives and we have to understand what level we are comfortable with of punishing innocent domain owners in an attempt to get at all the guilty domain owners. Both of those, to me, are important.

LG FORSBERG: Thank you, Reg.

I will now turn to Joanna, our At-Large end user representative. And I would like to ask you: In what way do you think that Reputation Block Lists providers could further help an end user that has been listed or has a problem with a reputation block list?

JOANNA KULESZA: Thank you, LG. I understand there are time constraints, but please let me briefly reflect on this rich discussion we've been having thus far.

This is Joanna Kulesza for the transcript record and the translation purposes.

Indeed, that is the core question the At-Large community is attempting to answer: How do Reputation Block Lists directly impact the situation of end users? Where can they turn to should their domain name been undiligent -- incorrectly qualified as harmful? So that to us is an issue of conveying the message from this session to the broader community but also representing end user issues and the development of criteria for Reputation Block Lists.

We welcome the opportunity to participate in discussions going on in the contracted party house and the GAC, as I was trying to just briefly put into the chat.

Usually, normally, end users would turn to their local service providers or law enforcement authorities to try and protect their domain, domain name, trademark, resources, services they are trying to offer. But these methods might prove ineffective because if an end user was to follow the thread of Reputation Block Lists, they would end up in this specific panel.

So to us, this issue is twofold. On one hand, we want to have a specific answer to our end users on how they can delist a website that was listed for the wrong reasons, that was listed inappropriately because, indeed, there is no malicious activity generated there.

Discussing automated and manual blacklisting, therefore, is to us an element of both building capacity among end users but also making sure that the criteria for blacklisting reflects their needs and expectations. While it seems as if at this point, as Peter noted in the Q&A pod, there is an elephant in the room that deals with governance and accountability of Reputation Block Lists.

As we could see from this discussion here today, various criteria is being taken into account. There is some conversation between those providing Reputation Block Lists; but, to us, to be able to understand how this entire system works, we would like common denominators or opportunities like this to better understand how the system works and be able to contribute to its development.

So, indeed, the criteria that will be used to blacklist certain websites is a discussion within the ICANN community where the At-Large voice needs to be heard.

We have heard about spam. The At-Large community has organized Webinars, sessions, internal policy discussions to identify DNS abuse categories to try and define these. Spam has proven to be one of the more controversial ones, being legal in some jurisdictions, illegal in others.

So looking at Reputation Block Lists, to us, an essential element is the criteria upon which these are set. We welcome the opportunity to participate in these discussions as well as the Monday's plenary on regulatory advancements, which will likely also impact the way that registries and registrars work in different jurisdictions.

I'm going to stop here. I know we're short of time, LG, but I believe this is a discussion to be carried forward.

Thank you.

LG FORSBERG: Thank you, Joanna. I would like to follow that up with a question for the Reputation Block List providers that ties directly into Joanna's statement.

How would you suggest that end users that have been listed in a reputation blocklist go about resolving that issue? Carel.

CAREL BITTER: Sure, yeah. So I don't know what others do. I can tell you what we do. We have a website where you can look up if -- you know, since we're talking about domains, if your domain is part of our datasets. If it is, in the majority of cases you can apply for removal directly, which will be handled directly, and the domain will be out of our datasets like at the next minute. We build the datasets every minute.

In certain cases, you will be forced to create a ticket because of, you know, certain metrics that we collect. And we have people working 24/7 to handle the tickets and will help you understand what the issue is, what we think needs to be done to get the de-listing, and a human will review the listing and then either help you understand what's going on or de-list the domain.

LG FORSBERG: Thank you, Carel. Ben, would you say this differs significantly in your case?

BEN COON: That's interesting. Our case where our list is mainly only used in, like, 24-hour segments because the phishing will go offline, but we do mostly the same thing that Carel was saying, where we will -- if somebody reports a false positive to us or if somebody reports something to us that is out of the norm, you know, we take a hard look at it, both not only what is going on in the automation of why that got picked up but also what's going on specifically with that.

You know, when we report something to registrars or hosting providers and they -- they push back on us, you know, we take that very seriously, and we look very directly into that to make sure that, you know, what we're putting out there is, you know, legitimately phishing or legitimately malicious. And if it's not, we remove it immediately. You know, like Carel was saying, within minutes.

LG FORSBERG: Thank you, Ben. It's clear from the discussion here today that there's not a generic feed -- feedback loop in place for report an indication of abuse as resolved or a false positive.

Roman, do you think that you would be willing to work on a standardized way of receiving information like this from a multitude of sources? And what would you -- what would need -- how would it need to be facilitated for you to trust it?

ROMAN HUESSY: I think that's an interesting discussion. I, of course, don't have a direct answer to that, but I think it would be a good way to -- to facilitate and to, yeah, generate trust in the community to have something like this. So that's definitely something that could be worked on. The question is just, yeah, who would be the umbrella for such a -- such a thing, whether this is an ICANN topic or something that must be solved by the industry.

LG FORSBERG: Thank you, Roman.

Joanna, I want to go back to you and ask a question. This is not directly targeted towards the Reputation Block Lists, but what fail-safes do you think are needed to make sure that registries and registrars do not act above their mandate when it comes to disabling domain names that have been indicated as participating in DNS abuse?

JOANNA KULESZA: I would support Matt's statement. This is Joanna, for the transcript record again. It's all about transparency. The more clear the criteria is and the more open the Reputation Block List providers are to discussions about the criteria that is being applied, the easier it will be for an average end user to understand the process and have their voice heard.

And I do firmly believe that this is where the At-Large can help to identify what these expectations are. The response could be regional. We have discussed this at Monday's session as well. The At-Large is working on a DNS abuse capacity building campaign that will start off regionally. So these responses could, indeed, be regional on behalf of the end users. The expectations could be different in different RALOs in different countries. We have seen what they look like in Europe during Monday's session.

But this would benefit, my understanding, the further development of the multistakeholder model.

So for us to be able to contribute to this discussion better, as Matt emphasized, transparency about these processes, understanding how they work and as the ICANN community has done forever, acting in good faith is something that I think will help us move forward to effectively advance Reputation Block Lists. Thank you.

LG FORSBERG: Thank you, Joanna.

I'm going to take an audience question in for the next one. This comes in from Volker Greimann, and he asks: Why do so many Reputation Block Lists not provide any evidence in the reports? Without evidence or specific reference that allows for verification of the report, it makes taking action exceedingly difficult.

And I will pass this on to Carel.

CAREL BITTER: Sure, yeah. So let's talk about the spam aspect of this. That's what I'm most familiar with.

We get our spam reports from either corporations with large ISPs where they share certain characteristics of the spam that they get with us or from our own spam traps or honeypots, if you will.

Spammers have become really good at embedding all sorts of little tracking details into URLs, into domains, into images, into the full emails. The scale at which we receive these things makes it basically impossible for us to make sure that we strip everything out of there and to make sure that any evidence that we provide has been cleaned of anything that identify our honeypots. So at that point a decision becomes if you can't share it without protecting your sources, then you can't share.

Like I said, on an individual case, we're always willing to work with giving people the evidence that they need, and if we can share it. But in an automated sense, this is just not a feasible thing to do.

To give you an idea, just one set of honeypots that we run receives between 2- and 3,000 emails per second. There's no way we could screen all of those to make sure that all the identifiers are out of there and to make sure that the data that we provide would be -- would be clean to not identify the sources.

So like I said, you know, if people who are -- who are doing remediation work need specific evidence, then the people who get our data for that purpose will know where to contact us, and they can get additional details, you know, where -- where possible.

LG FORSBERG: Thank you, Carel. And sadly, that ends today's session as we running out of time.

Looking at the answers to the poll questions from earlier this session, I would say that our audience has been quite broad.

There are people that have not heard of a lot of different types of abuse and there are definitely people that have heard. There are people that haven't had to deal with abuse mitigation and there are those who have.

Hopefully some of us have learned quite a bit from today's discussion and can bring this forward in their work with this particular subject.

I would like to thank all of the panelists for their time and participation, the ICANN staff for their invaluable support, and last but not least, the audience that took the time out of their day during this busy meeting week to sit down and listen.

Understanding Reputation Block Lists, their function, purpose and how going about using them or contacting them to correct an error is a very important step in developing ways of working with them and against DNS abuse.

There is a long road ahead with this, but I personally think we made some headway today.

Again, thank you all for participating.

BRENDA BREWER: And we do have a final polling question. We will open that poll question up momentarily. Please stand by.

And the polling question should be on your screen. Do you know how to report a security threat or take other steps to mitigate those threats? Please answer yes or no.

Again, the question on your screen is: Do you know how to report a security threat or take other steps to mitigate those threats? Please reply yes or no.

Thank you. And I will close the poll and share the results.

Thank you very much for your participation today. The meeting is now adjourned.

Thanks, everyone.

BRENDA BREWER: You may stop the recording.

Thank you, everyone.

[END OF TRANSCRIPT]