ICANN73 | Virtual Community Forum - Plenary Session: Evolving the DNS Abuse Conversation
Wednesday, March 9, 2022 - 10:30 to 12:00 AST

BRENDA BREWER: May we have the recording --

Recording in progress.

BRENDA BREWER: Hello and welcome to ICANN73 plenary session: Evolving the DNS abuse Conversation.

My name is Brenda Brewer, and I am the remote participation manager. Please note that this session is being recorded and follows the ICANN Expected Standards of Behavior. To ensure transparency of participation in ICANN's multistakeholder model, we ask that you sign into Zoom sessions using your full name. For example, first name and last name, or surname. you may be removed from the session if you do not sign in using your full name.

Interpretation for this session includes Arabic, Chinese, French, Russian, and Spanish. Click the "interpretation" icon in the Zoom toolbar to select your preferred language output.

During this session, questions or comments submitted the chat will only be read aloud if put in the proper format as I will note in the chat. I will read questions and comments allowed during the designated times for this session. During the community session portion, if you wish to speak, please click "raise hand" in the Zoom toolbar.

Before speaking, please mute all devices and notifications. Also, ensure that you have selected your preferred language input. Please speak clearly and at a reasonable pace to allow for accurate interpretation.

Once the session facilitator calls your name, kindly unmute your microphone, then state your name for the record. To view the real-time transcription, click "closed captioning" in the Zoom toolbar. And now please welcome session moderator, Graeme Bunton.

Please begin.

GRAEME BUNTON: Thank you, Brenda. Good morning, afternoon, evening, everyone. Thank you for joining us for today's plenary: Evolving the DNS Abuse Conversation, Malicious Registrations and Compromised Domain Names.

My name is Graeme Bunton. I'm the executive director of the DNS Abuse Institute.

The first thing I want to do is apologize for the length of the introduction I'm about to give. We have got a complex topic, some amazing panelists, but a fair amount of groundwork that we need to lay before we really get into it. So please bear with me.

Our first presenter, Maciej, is going to give us a far more robust introduction to the substance of the topic. But I want to make sure for the rest of my intro here, everybody is on the same page.

So here's my attempt at a two-sentence explanation of what we're talking about today. In simple language, our session today is about exploring the differences between the mitigation processes for domains that were deliberately registered to do harm and domains that are used by websites that have been compromised, or hacked, that are also being used to do harms. So in one case, someone is deliberately trying to do bad things with a domain name and in the other they're not. And we need to understand how that complexity lands within our ecosystem.

So before we get into -- really into the substance, I want to go over a few things about the goals of the session and the scope of the session and give you a bit of an outline for how I think we're going to structure this conversation.

Brenda, if I can get the second slide, I think, please.

Plenary goals, so this is what we're going to try to accomplish today in the short 90 minutes that we've got. We want to try and develop a community understanding of why this distinction between malicious registrations and compromised websites is important. We want to talk a little bit about how that distinction is to be made but sort of a technical problem in some senses, so we're not going to dive too deeply into that.

But then we really want to dig into -- and this is the meat of this conversation. We want to develop a community understanding of what could be done in either scenario of malicious versus compromised and dig into that.

And then, lastly, we want to talk about potential activities, what could we do about this problem, who should be doing it, what's the role of the community.

So let's go on now to the scope. And this is important. I have been doing ICANN stuff for too long now and have been involved in lots of plenaries. And I think there's a relationship between how specific we are in a topic and how much value we get out of sessions like these. And so I really want everybody here to join me in being quite narrow in our scope today.

So there's some assumptions we're making to have this conversation. And really what we're talking about here is that a registry or registrar has received a report of abuse. How that report got in the door really isn't up for discussion today. We're going to assume that the abuse has been verified. There is something that is for real happening. We don't need to discuss whether that is happening or not. So that verification is sort of out of scope.

And we're primarily focused on examples of malware and phishing today. That may not be how you define DNS abuse. It is probably the center of most people's definitions of DNS abuse, but we're not here to discuss that definition. Let me be very clear. We are not going to engage in a definitional discussion of what is or is not inside of DNS abuse.

We're really here to understand this distinction between malicious registrations and compromised.

So what we're going to do going forward for the rest of this conversation is really keep it constrained to this topic. We will try and capture out-of-scope comments and questions to see if there's another place for those questions in the future.

But I'm going to be pretty robust and vigorous in keeping us on topic. And so I ask all of you, when you're submitting questions to the Q&A pod or discussing in the chat to keep us on topic and please be understanding when I'm trying to do that for all of us as well.

Next slide, please.

So this is a bit of a map of the conversation that we're going to have today. So we've got a report of DNS abuse at a registrar/registry. We need to talk a little bit about the "why," should we make a distinction here. Is there -- should it really just be a generic abuse process for all types of harms? If we don't agree with that and we do need to make a distinction here, we need to talk a little bit about how that distinction is made. You know, what are the attributes of registrations that are going to lead us to making that choice?

And then we have two streams here. We've got a process that we should talk about for maliciously registered domain names, domain names that are deliberately engaged in online harms.

And we have another process for compromised websites and what we should be doing there.

And a bit of a spoiler, I think the real meat of this conversation, the real thorny bits that we need to figure out as a community is what we do on the process for compromised websites. But we'll be diving into both of those as we go forward.

So we've got a delightful panel here today. We're going to start with an introduction from Maciej, professor at the University of Grenoble -- gre-nobe, I think, if you are French -- who is really going to give us some data and an understanding of the lay of the land. We're going to take about 15 minutes to do that.

And then we're going to go right into a panel discussion. And our panel today has got some very illustrious people on it. Thank you very much to all of your for joining. We have got Lori Schulman joining us from the IPC. We Chris Lewis-Evans from the GAC Public Safety Working Group; Alan Woods from the registries; Reg Levy from the registrars; and Rod Rasmussen from the SSAC. And so we're going to dig into these topics as we see them before us.

Roughly broken into two sections. One, we'll do the introduction, and then we'll do a bit on the malicious and then we'll do some more conversation on compromised websites.

We're saving a bunch of time at the end for audience questions. I will say the Q&A pod will be open for this entire conversation, and I encourage you to use it with the caveat that we are really trying to stay focused on the topic at hand today.

And so if your question doesn't get answered know we'll capture it for later potential discussions, but please don't be too upset if it doesn't get answered in the context of what we're discussing here today.

So I think that's all I wanted to share for a slightly long-winded introduction, but I think hopefully we've got some expectations, we've got some clear goals, some clear scope, and maybe it's time to dig right into it.

So with that, I think I'm ready to pass it over to Maciej who is going to give us a real "what's happening today" in this day of malicious and compromised domain names.

So Maciej, go ahead.

MACIEJ KORCZYNSKI: Thank you, Graeme, for the introduction.

Hello, everyone. Today I will discuss briefly the problem of compromised versus maliciously registered domain names. And it will be mainly based on the COMAR project, a research project funded by AFNIC and SIDN, two registrar operators of .FR and .NL domains and are also based on the technical part of the study on DNS abuse commissioned by the European Commission.

Next one, please.

So here on the top we can see a URL that was blacklisted by a Phish Tank, and below you can see malicious website, the screenshot of malicious website. So the question we are trying to answer today is the domain name maliciously registered. Next one, please.

So to answer this question we need to investigate this case a little bit more. So when we visit the registered domain name, then we can see that there is no meaningful content, and also when taking a look at the WHOIS information, then we see that it was registered just two days before the actual URL was blacklisted.

Next one, please.

So that provides us quite strong evidence that the domain name is maliciously registered and abused to serve illegal and abusive content, phishing of credentials and trademark infringement.

So what are the implications? What intermediary should mitigate abuse from the technical perspective? So it should be DNS service operator -- for example, registrar, eventually TLD registry -- and also the hosting provider. And why this is so important? Because if we only suspend the domain name and not the malicious hosting, then the attacker might simply might simply register another domain name and point it to the malicious hosting. On the other hand, if we only suspend the hosting but not without blocking the domain name, the domain name could be reused by the attacker in other attacks, in other phishing campaigns.

So to increase barriers for abuse and economic cost to the attackers, the mitigation should be both at the DNS and hosting level from the technical perspective.

Next one, please.

So let's investigate another case. So here we have another domain name -- sorry, another URL, malicious URL, blacklisted, if I remember cell, by Anti-Phishing Working Group. And below we can see the screenshot of the malicious web page. So the same question here, is the domain name maliciously registered?

Next one, please.

So when we visit the registered domain name, there is a website with a legitimate content, and the content also corresponds to the domain name itself. When we take a look at the WHOIS information, the domain name was registered back in 2014, so most probably the domain name itself is legitimate.

Also, when we take a look at the malicious URL, then we can notice wp-includes, which indicates that the actual website has been hacked by exploiting vulnerable Word Press installation.

Next one, please.

So the domain name is legitimate, but the website was compromised and abused to serve illegal and abusive content, phishing of credentials, and trademark infringement.

So what are the implications from the technical perspective?

So generally should not be a registrar or TLD registry that blocks the domain name because that might cause collateral damage to the registrar, to the business behind, and also legitimate visitors of the website. So the mitigation should be at the hosting level by hosting provider or the owner or administrator of the website. And there are two things to do. First, to clean the malicious content, and also to patch vulnerable and Word Press installation.

And now, who should do it? It should be either the hosting provider if the hosting is managed, for example, on the shared hosting platform where the hosting platform controls all the software, including vulnerable software, or the web admin if we're dealing with unmanaged hosting and its administrator who controls vulnerable -- vulnerable software in this case.

Next one, please.

So how legitimate domains are abused? So from our analysis, we find that the domains are mainly abused at the website level. So there is a vulnerable software that get exploited, for example, content and advanced system. Sometimes it happens at the DNS level. An example here could be domain shadowing where the attackers first will try to phish for credentials of registrars, registrants to get to their registration panel. And once they log into the registration panel, they can add, for example, subdomains that could be used in, let's say, phishing attacks.

Next one, please.

So what are the existing approaches to distinguish between legitimate but compromised domain names from maliciously registered domain names?

So there are two approaches. So the first approach is based on heuristics, and this is very often used in industry reports. So one of the heuristics is the domain name age. So as previously illustrated, the time between the registration and blacklisting; registration in bulk patterns; and also patterns in registered domain names. So, for example, for phishing attacks, brand names or misspelled versions of the targeted service, such as PayPal.

And second group of approaches are machine-learning methods. And here example is COMAR developed in Grenoble Alps University, funded by AFNIC and SIDN, and it's a fully automated approach, and the idea is we collect data related to hosting to the website, to the structure of the URL, we determine specific QRs in domain names, and so on and so on, and we extract the 38 features. The project is fully automated based on modeling, and we show that its accuracy is up to 97%.

Next one, please.

So what is the relation between the type of abuse and compromised versus maliciously registered domain names? So next one, please.

So here we show the distribution of compromised in blue and maliciously registered in red, domain names per abuse type. For spam and bought net controlled, generally, the attackers need to control DNS. However, for phishing and mall ware, this is not the case. The attacks can be launched using maliciously registered domain names or compromised websites, also free services.

And we are about 25% of phishing domain names and 41% of mall ware distribution domain names are registered by legitimate users but most probably compromised at the website level.

Next one, please.

So what is the variation across different TLD types? So in the second part of 2021 we observed that for new gTLDs, almost 98% of the domains were labeled as maliciously registered. If we take a look at European Union ccTLDs on the left, we see that 42% of the domains were labeled as compromised websites. And what would be pos- -- what are the possible reasons for that? We suspect that in European Union and ccTLDs we will have less speculative registrations, more fully featured websites serving meaningful content, meaning also having deployed different software that could be potentially vulnerable and exploited at scale by cyber criminals.

Next one, please.

So here we see the variation across different TLDs.

Sorry, I think we lost the connection. I do not see the presentation.

GRAEME BUNTON: You're back now where we can hear you, Maciej.

MACIEJ KORCZYNSKI: Okay. But I cannot see the presentation.

GRAEME BUNTON: It's slide 13 that's still up.

MACIEJ KORCZYNSKI: Yeah, okay. I'll just switch to a -- to my local version. Okay.

So here we can see the total number of domain names, abuse domains per different TLDs and number of maliciously registered, number of compromised and in particular, in the last column, we can see the percentage of the malicious registrations to all the domains abusing particular TLDs. Next one, please.

So here for certain TLDs we see the percentage of maliciously registered domain names almost 100%.

Just to mention .TK and .ML domain names were Freenom gives domains for free to their users, and this is an attractive back door for, let's say, phishers.

Next one, please.

On the other hand, we see also TLDs with. For example, .BR, 34% only of maliciously registered domain names. And this can be driven by multiple factors. I would like to pay your attention that those results needs to be taken with caution because of the limitations of the classifiers but also because of the limitations of the blacklist themselves. They might not represent the entire cyber space.

And why is that? Because some blacklist providers might concentrate on the maliciously registered domain names using this -- by identifying certain keywords and so on.

So one thing also I would like to mention here is that what we see in blacklist might be different also what the TLD registries receive in their help desks when they're analyzing the complaints from the victims of the attacks once more because those blacklists might not be representative enough.

And my final remark is that from time to time we see quite big variations for the same TLDs, for example, .INFO or .COM, from one month to another. And what could be the reason? One possible reason would be that one of the resellers, for example, gives big discounts to the domain names that are exploited by the attackers. And we would, in this case, see increase in the percentage of maliciously registered domain names to the total number of abuse domains. Or, for example, from time to time, we see some vulnerabilities that are being discovered, for example, in content management systems, affecting hundreds of thousands of domain names. And this is also a low-hanging fruit for the attackers. They might exploit them at scale. And then, in this case, within the shorter period of time, we would see the decrease in the percentage of maliciously registered domain names to the total number of the domains.

Next one, please.

So thank you very much for your attention. I hope it will help in starting the discussion. Thank you.

GRAEME BUNTON: Thank you, Maciej. As you were going through that, you answered a number of questions that were coming in, so I really appreciate it.

There has been a little bit in the chat, though, Maciej, around can we do this with machine learning? Are there algorithms for detection and their accuracy? And I don't think we want to go too deep on the "how" to do this, but I think it would be useful if you could give us maybe two or three minutes on the various attributes of domain names themselves, or the websites, and how tools might automatically try and do this or, you know, what tools are useful for humans to try and do this where they don't have this technology in place. Could you do that for me?

MACIEJ KORCZYNSKI: Absolutely. So regarding the first question, we did the analysis of the importance of the features in the COMAR classifier.

So the COMAR classifier, as mentioned during the presentation, is a fully automated machine learning algorithm. And features that are very useful are content-based features. So, for example, if we detect on the website that there are a lot of different technologies that are used to build the website, that is an indication that the website was compromised.

If we see some assertive keywords such as "PayPal," but not only brand things but very specific keywords, like a verification, we also did such analysis of keywords. Then it's an indication of the website being maliciously registered.

And, you know, some things can be recaptured by the classifier like the number of technologies, which is much harder for humans.

On the other hand, if we want to do it manually, then a couple of things like we've seen in the examples, the time between registration and the blacklisting. If it's really short, then it's a strong indication that there is -- that the domain name itself is maliciously registered.

If we see no meaningful content on the registered domain name or the actual, let's say, phishing website on the registered domain name, this is another indicator that the domain name is maliciously registered.

And those two examples are good examples where both humans and also our machine learning -- machine-learning approach can detect those things.

GRAEME BUNTON: That's great. Thanks, Maciej. So there's quite a bit going on in chat. I can't follow the chat because it's too much. So if you do have a question that you would like answered, and, again, is on topic, please use the Q&A pod so we can try and capture that.

So there are a number of questions more on the "how" of this detection. A couple things you might look at is the DNSAI produced a best practice on this relatively recently.

I'm sure one of my friends can go find a link and dump that in the chat. Please and thank you.

And I also did, only me, a presentation on architecting these sorts of solutions on Monday as part of Tech Day, which I think provides a framework for how to do this.

But I think we need to move on a little bit from the "how." And it's clear there's lots of interest in this technology or human approaches in making these determinations about whether a registration is malicious or not. But I think we need to move on into the "why" and then the "what."

And so this is going to be my first question to the panel. Panel heads up.

There's this gating question here, is: Should we doing this at all? Does it make sense to bifurcate our mitigation processes? Do we need to distinguish between malicious or compromised or not? You know, should we do the same approach for everything, a generic abuse approach; or should we have separate processes?

And so I'm curious if anyone disagrees with the premise to start. Do we treat all abuse as the same?

And so let's start there and see if anybody from our panel wants to weigh in on that one quickly.

I will pick on you. Reg, please.

REG LEVY: Sure. Yeah. I think it is absolutely necessary to make this distinction. We have a lot of customers who use commercial website creators that require regular updates. And if they do not perform those updates, they are often vulnerable to compromise.

So reaching out to the registrant to discuss with them the fact that they need to do something on the domain name that they may have not ten years ago and haven't thought about five years since they put up the website. And they are just using email and they assume people can go to the website, see their information, and be able to contact them is a process. And we need to make sure that their business is not impacted because there was a compromise on the domain.

GRAEME BUNTON: Thank you, Reg.

I have got Alan and then Chris and then Lori. And then my hunch is we don't need too deep on the "why." We'll see. And then we'll move on a little bit.

Alan, please.

ALAN WOODS: Thanks very much. Alan Woods for the record.

So, yeah, I mean, I think it is important that we do make this distinction. And I know it's the reason that has been trotted out a lot through many, many years. But that is, from our point of view as a registry operator, when we take action against a domain, there is a lot of collateral damage. Therefore, we do not want to victimize another victim, and that is the registrant, in an instance of compromise.

So I think we need to be very clear that if we are going to take action, that we should have an idea and an appreciation of the difference. I was going to a "nuance," but it is a difference between those two different types of victims as well.

GRAEME BUNTON: Yep. Thank you, Alan.

Chris.

CHRIS LEWIS-EVANS: Yeah, thanks, Graeme. Chris Lewis-Evans for the record.

Just to agree here, I think we should be treating these differently. As Alan said there, we have two different sorts of harm being caused here. So in a maliciously registered domain, there is primary harm. And then in a compromised domain, we have a primary and then a collateral harm being caused.

So with the compromise, we also have two types of victims, the primary victim and then those that may be affected by sort of collateral harms. So really we need to be able to treat them both and provide both with sufficient help.

GRAEME BUNTON: Thank you, Chris.

Lori, please.

LORI SCHULMAN: Yes. I was going to say that the IPC agrees that there is absolutely a need to distinguish. It's important to distinguish between malicious and compromised domains at least at the front end in terms of how quickly we can respond to a particular issue.

But at the same time, though, I wouldn't want to be lost on this issue of distinguishing who is the actual victim because there is the end user, the end user that's being phished or the end user that is the subject of a malware attack.

But there is also, particularly in the case of small businesses, potentially on that compromised owner having their business compromised, their reputation compromised. And I don't think we should make the assumption that the registrant who's running the business -- or running any -- doesn't necessarily have to be a business per se, running any entity off their website, wouldn't prefer to have their site down for a certain amount of time if it, in fact, is protecting their customers or their reputation.

GRAEME BUNTON: Thank you, Lori.

And so where I got to from that is it seems, at least within this panel, that no one disagrees with this question, that we really should be making this distinction. And that's great.

But it also makes our lives more complicated. And so now we need to begin to dig in a little bit on what that means. Having said that, we've got a number of questions in the Q&A pod that I think maybe we should try and address before we move on too much so that they stay relevant. I might try and pick some of these and assign them to panelists. And we'll try that out.

There was a question from Greg Shatan that I think belongs to Maciej. And so he was asking: How does abuse at the mail server level fit in your proposed approach?

And abuse at the mail server level I think is pretty interesting. I don't know enough about it.

Maciej, do you have any thoughts on how to tackle that?

MACIEJ KORCZYNSKI: The mail server?

GRAEME BUNTON: So I think Greg is asking about where it might be phishing or malware via email.

MACIEJ KORCZYNSKI: Would it be possible that the author of the question simply ask the question?

Perhaps it will also be easier to clarify the question.

GRAEME BUNTON: I'm a little bit cautious about doing live voice here on this panel in this method.

Maybe Greg can elaborate a little bit in the chat, and we'll come back to it.

MACIEJ KORCZYNSKI: Thank you.

GRAEME BUNTON: And see if he can clarify.

MACIEJ KORCZYNSKI: Thank you so much.

GRAEME BUNTON: What else can we try and answer before we move too far ahead. Oh, boy, there's lots in here. So bear with me while I try and pick some that are going to be relevant.

Another question for Maciej was from Samaneh. She was wondering if the ML method and the feature used in COMAR also includes some of the heuristics you pointed out in the first method. If so, which ones as examples?

MACIEJ KORCZYNSKI: Thanks for the question. Yes, so we included almost all the features that are used in heuristics in those methods apart from bulk registration. This is the only feature we did not include. And the reason is -- I would say that there are two reasons. One reason is that we -- the main reason is that simply the COMAR method should distinguish maliciously registered versus compromised websites only based on the data collected and related to one specific case without getting -- getting information from, for example, other maliciously registered domain names.

But apart from this one, all the heuristics are implemented in the fully automated COMAR system.

GRAEME BUNTON: Thanks, Maciej.

I think there's another one here for you from Michael Palage about he would be interested in your opinion on the high percentage of compromised domain names within European ccTLDs and whether that would -- would this not likely be attributed to a growing number of ccTLDs doing identity checks? Is it easier for malicious individuals to compromise a domain/website as opposed to register a domain name with fraudulent or synthetic registrant data?

MACIEJ KORCZYNSKI: If I understood the question correctly, the question is why we see less maliciously registered and more compromised websites with European Union ccTLDs. Is that correct?

So the answer I tried to discuss it a little bit during the presentation. But we can only speculate because we did not make any measurements.

But I would say that in a European Union ccTLDs, we have, as mentioned -- as mentioned before, many less park domains, not that many domains that are speculative. There are websites behind the domain names. So if there are websites, their users also take care of them. They put -- they deploy different software. And because we see a lot of different software on the website, some of the software might be simply -- might be simply -- might be simply exploited.

And I think the second part of the question was, like, why we see less maliciously registered domain names. This is also purely speculative. But there are many initiatives at ccTLDs preventing malicious registrations. Just to mention that at the E.U. ccTLD, there is a system similar to Premadoma that detects the domains -- maliciously registered domains at the time of registration. Or there are other ccTLDs that actively fight and try to prevent those malicious registrations, like SIDN AFNIC, for example.

But this is more from my experience and research and practices that I see at the ccTLDs. It doesn't mean that, for example, other ccTLDs, other groups of ccTLDs, do not deploy those methods.

GRAEME BUNTON: Thanks, Maciej.

So, boy, we've got a lot of questions and a lot going on in the chat. We'll do our best to stay on top of all of that.

Panelists, if you feel like answering a question either live or via answering in the Q&A pod, please feel free to do so.

But I think maybe from this point, we should move on to a little bit of that left side of that diagram and begin talking about what the process might look like for a malicious registration and considerations around that. So just to make sure, again, we're all on the same page, we decided we should distinguish. We talked a little bit about how we might do that decision. So we've got some attributes of the domain name. Maybe it's ML. Maybe it's a person doing it. Now we've got to figure out what we're going to do.

And so at the registry and registrar level, there's not a ton of options but perhaps we should explore them a little bit.

And this time, I might turn to Rod. Rod, if you have thoughts on what activities a registry or registrar should be doing when they have encountered DNS abuse that they believe to be maliciously registered.

Do you have thoughts there?

ROD RASMUSSEN: Yes, thank you. This is Rod Rasmussen.

So once you've determined -- you made a determination for whatever methodology you've used and put that aside because we're not focused on. We've said, okay, we have decided this a malicious registration, what can I do?

And we have -- as a registry/registrar, you have really very few options. We're going to mainly have the same effect which is to remove the domain itself from the global DNS.

And there's a few ways you can actually do that. You can delete it right on the spot and say basically we're going to remove this registration. If you have a malicious registration that's been done, you know, within the last few days, you actually have the -- take the advantage of being able to get some sort of a financial recompense for that. In other words, you get your money back. So there's an advantage of it, but there's a disadvantage of doing it in that whoever registered it initially can turn around and register it again either using the same registrar or somewhere else and just re-energize their scheme, whatever it may be, even assuming you -- at the same time somebody is mitigating the malicious content wherever it may be, they can obviously -- you know, there's lots of compromised sites that can reestablish that. So that has advantages and disadvantages.

You can suspend the domain. In other words, you do not delete it, but you actually put it into a suspension status. That would remove it from the DNS, and then it will, for whatever the lifetime of that registration is, will be sitting around in a suspended status. And that's just something you then have to manage as the registrar or registry or both.

Then you can -- there's other active mitigation measures you can take as well. For many years, the Anti-Phishing Work Group has provided a landing page for phishing sites, so you can actually redirect the -- if it was a phishing attack, for example, you could redirect -- you could actually change the DNS and point it at that phishing landing page. Other people have done similar types of things. If it's something like malware you could create what's called a sinkhole and create the ability to inform victims that their machine's been compromised. This has been done in many different ways either directly by providers, by security companies, by law enforcement, et cetera. So you can actually take an active role in trying to let various victims of malware know that their machines are infected. And in order to do that you may need to transfer the domain name to another entity, either from a registry -- registrar -- a registrant perspective, excuse me. So, for example, the FBI has seized domains and had them transferred. Microsoft has done this many times as well.

You can also -- There's the -- utilizes the register of last -- Registrar of Last Resort which has been set up to take in command-and-control domains for malware and then provide that data back on an automated basis to victims. So there's several different options there, but that does involve some work in making sure you have a process and legal paperwork in place.

And one last thing I'd like to add to this is once you have determined that you have a malicious registration of some sort, it's probably a very good idea to look and see if there are other domains that are lined up or being used by potentially the same registrant or registrant account. In taking a look at that account, it's really important -- and I think Reg or others might be able to talk to this a little bit further -- is to understand whether that account was set up by the actor or was compromised and has had domains added to it because they may be a victim of phishing themselves or some sort of credential theft.

And then the other thing to do is look for a pattern of a series of accounts that may have been created under different aliases, et cetera, to look for abuse across a wider swath of domains and maybe registrant accounts that could lead to -- would typically support a large-scale campaign, because some of these actors are very clever about how to hide themselves from diligent registrars trying to keep them off their systems.

So there's some thoughts for you.

GRAEME BUNTON: Thanks, Rod. Boy, that was great. And there's a lot in there.

If I can recap briefly, registrars sort of have three or -- registrars more so than registries, although registries could participate in this, but you can delete, which probably isn't great, you can suspend, or you can point or redirect. And there's going to be reasons to do all of those three. And then you're going to want to check the account, and you're going to want to see if that account fits a separate pattern. And I think all of those are pretty useful inputs for people who are mitigating abuse.

I'd be curious from Reg or Alan on our panel how they see those things. Do they employ those methods frequently? If so, why? Why not?

But before -- while you guys think about that, Lori, please go ahead.

LORI SCHULMAN: Thank you. I have a question to Rod's point about looking at patterns of abuse and looking for a wider -- wider swath of examples. And so today, with the constrictions that we have with privacy legislation and current policies, is the research now limited only within a particular registrar or can the registries look across a broad scope of registrars where it might make more sense for a registry to do this investigation rather than a registrar?

GRAEME BUNTON: It could well. Thanks, Lori.

Rod?

ROD RASMUSSEN: Yeah, let me respond to that. Yes to both (laughing). The -- And the -- the -- So a registrar has the unique capability of being able to see the things that have been redacted from the public view when it comes to contact information, et cetera, which was a valuable asset, a very valuable asset in doing these, the kind of heuristics, et cetera Maciej was talking about earlier. The -- But the -- And they have the ability to see where signups are coming from, credit cards being used, all that kind of thing. So internally, they have a lot more data to be able to look at things.

Externally, a registry can actually have a really good idea that there's a pattern going on, particularly if it's doing something like supporting what's called a domain generation algorithm. Those are used by malware where there's a predetermined set of domains that would be registered in order to supply command and control for malware families. If you see a series of those across different registrars, you know there's a pattern there.

They can also take a look at things like DNS hosting and the way the domain is actually configured on the Internet itself. So if you take a look at either the way that they're set up for particular DNS servers, nameservers, or eventual hosting IP addresses, you can -- you can often detect these kinds of things, at least it's suspicious. It's something you might want to take a deeper look into and ask registrars about, to take a look at those accounts and see if they're legitimate.

GRAEME BUNTON: Thanks, Rod.

And just a thought, maybe, for Chris. If you -- and you don't have to -- I'm going to go to Reg and Alan before I get back to you, but before we move on from this topic and compromised, I'm curious if, from a law enforcement perspective, on the malicious registration whether you're generally happier for just disruption or whether that -- for malicious registration you're often engaged in something like an investigation where you want more out of that process.

But Reg and Alan, briefly. Then maybe Chris. And then I think we're going to try to move on to the compromise side of things, and then we'll get to more Q&A.

Reg.

REG LEVY: Thanks, Graeme. This is Reg Levy from Tucows. We're a wholesale registrar, so our way of approaching this is going to be slightly different from a registrar that has a direct relationship with the registrants.

We work primarily with our resellers, and we look for patterns based on reseller. So when a lot of abuse starts coming from a particular reseller, we reach out to them and say, hey, we're experts on dealing with DNS abuse. How can we help you? And typically we resolve the issues on the basis of that.

That said, we also have an in-house reseller, and when we see a spate of maliciously registered domains coming from them, we can take action directly against the registrant in question as long as it's an identifiable party.

As has been mentioned in the chat, very frequently fraudsters don't use the right -- their own names or the same names to register a bunch of domains. So looking for patterns like that is not always useful.

That said, and I think the initial question was do we look at when the domain was registered in order to make a determination about whether it is malicious or compromised? And yes, we absolutely do. Unfortunately, a lot of AIs, which is being mentioned as well in the chat, misses things or provides -- or creates too broad of a spectrum for takedowns. So we need to do a lot of manual checking of what the AI actually presents to us.

One of my favorite stories is actually involving a domain-generation algorithm. Some poor idiot managed to register something that was exactly the random string of letters that a DGA was using, and it was their long acronym for a very small women's soccer team in -- I forget exactly where in Central America. So (laughing) we worked with law enforcement and the registry to allow them to actually register it.

But, again, it is an example of these sorts of computer-based algorithms are going to cast too wide of a net. Sorry.

GRAEME BUNTON: Thank you, Reg. And an enjoyable anecdote.

Alan and then Chris. If you don't mind keeping it brief so we can keep moving on. Thanks.

ALAN WOODS: Sure. I suppose very quickly. It's great to go after the registrar because that's pretty much the way the registry should be going in this particular conversation. And that is I am everything that Reg said, just a level up. Where she's talking to her resellers or talking to her registrants, I'm talking in the first instance to my registrar first. And I think it's very important that from a registry point of view, when I'm in that position, yes, we will look at those indicators, and my back-end platform TDNS will do this for me to an extent, but it's about evidence-based escalations to my registrars so that they can have much more targeted conversations with their registrants.

And I think at the moment we've kind of slipped into the detection conversation, which I want to move away from because we're talking about compromised here and how -- what do we do and how do we identify compromise. So I think it's an important aspect on that one.

The one thing I do want to say, and I think it's important as well, is from a registry point of view. And just I wanted to push back slightly on something Lori said earlier about I think small businesses would prefer or wouldn't mind if their business went down. I think a lot of small businesses would absolutely disagree with you on that one. And I think we need to be very clear it's about the proportionality of the harm. Is the harm being caused to registrant or to end user more? And if we can work on that proportionality basis as opposed to making somewhat larger statements or broader statements, I think it's definitely about proportionality, and we need to talk about here.

GRAEME BUNTON: Thank you, Alan. And that is going to be a very good segue into the next component of our discussion, and I'll come back to that balance of harms in a brief moment after we hear from Chris.

CHRIS LEWIS-EVANS: Yeah, thank you. I'm Chris Lewis-Evans, for the record.

So sticking with the maliciously registered domains, Graeme, which I think is where your question for me, yes, at the end of the day we're talking about malware and phishing here, so generally there will be some form of investigation. And it will come as no surprise to anyone that generally they don't register one domain.

So the -- the domain will be under investigation, and any work that we can do with registries, registrars, or hosting providers to identify further domains, to switch it from a reactive to a proactive piece of work is always going to cause further victims and further harm being caused. So by all means, yes, that it is always something law enforcement are looking to do, is to switch up and prevent harm.

Thank you.

GRAEME BUNTON: Thanks, Chris.

Okay. So we've got about 35 minutes left, and we want to make sure that we get to a bunch of these questions in the queue, and there's a lot of activity in the chat. Thank you all very much right now for that great engagement that we're seeing. But we have, I think, so far done the easy work, which is we figured out we want to do this distinction. We talked a little bit about how. We talked about some great stuff from rod on what to do in the circumstance of a malicious registration. And now we have the other side, which is where I think it gets more complicated which is where there is an example of DNS abuse, and our examples say malware and phishing on a compromised website. So the website itself is benign. Might be a small business, might be someone's blog. Doesn't really -- or maybe it does, and that's the conversation we're about to have.

What is the process for figuring out what to do in this circumstance? What is the balance of harms that people are typically doing?

And so maybe I'll start with sort of a provocative question for Reg and Alan again, which is going to be if you've determined that a website is compromised and it's serving phishing or malware, it's engaged in DNS abuse, is there a circumstance where you would take it offline? You would suspend the domain, let's say, for a compromised website.

Alan?

ALAN WOODS: Thank you. It's almost as if my magic, it's about disproportionality of harms. So that is an interesting question and I think we need to be very clear as a registry that we maintain a very, very blunt instrument and that is the taking down of everything on that website, everything on that domain, any emails associated with that domain. So in order for the registry to actually say do you know what? In this instance I haven't gotten any response from the registrar, and I have gotten no response from the registrant, and harm which is objectively exceptionally large is still occurring. Yes, I mean, there is always a point that we might take down that. but it is on proportionality of harms. Are we talking about a -- sorry, I will slow down for the interpreters. I get excited.

We are talking about, you know, things that are harm to human life. We're talking about things such as child sexual abuse material, things which are exceptionally, exceptionally derisive and divisive. And we need to be very clear that the registry should not be the point of taking down where at all possible, but where it's necessary, we have that option.

GRAEME BUNTON: Thanks, Alan.

REG LEVY: And to follow on --

GRAEME BUNTON: Go ahead, Reg.

REG LEVY: To follow on to that, we tend to use suspension as an option of last resort for compromised domains as well. We will reach out to the reseller and to the registrant directly and say, hey, something's up. Can you fix it? And based on their response or lack of response, we may start resetting various records one by one so that we can turn off mail, we can -- if that's where the compromise is happening. We can reset a nameserver if that's where compromise is happening.

Sometimes that will prompt people to respond and say, hey, I didn't realize that this email from you was something I needed to respond to, but can you put my (indiscernible) online and then I can fix it. Which is another piece of it, that sometimes the domain -- we need to allow the domain to resolve in order for them to log into their website and fix the issue.

So suspending the domain will get rid of the malicious use, but it won't allow any other use and it won't allow mitigation of the abuse.

GRAEME BUNTON: Thanks, Reg, thanks Alan, for that input.

I'm curious from, say, Lori, Rod, Chris, if there's information that you think should be included in that balance of harms test that registrars are often engaged in that you think is maybe underrepresented or could be -- should be elevated in those thoughts as people are thinking through what to do with a compromised website that's engaged in DNS abuse.

LORI SCHULMAN: Yes, Graeme, I'm going to, if you don't mind, take that first because of what's going on in the chat. Because I made a statement in the chat that has definitely enlarged the conversation. But this goes back to something that Reg is explaining, which I agree with, that when you're looking compromised domains, you're not looking necessarily at a set of facts always. And that was the point I was trying to make in the chat, but I think it got a little overinterpreted, right?

When you're looking at compromised domain names, you're looking at a completely different decision path in terms of what you're going to do, when you're going to do it and how you're going to do it because there is the involvement of a -- presumably an (indiscernible) pattern, a live website that's operating, that's offering a service or some other benefit that if we were to suspend the domain, then that benefit goes away. And that could include somebody's income. Very well noted.

But the point that I'm trying to make is depending -- and this goes to Rod's point about -- and to Alan's point about what are we talking about, you know, in terms of the phishing or the malware. You know, has the site been compromised and now it's running pornography? And I'm not even going to say any -- I'm going to step back and say child pornography, so we don't get into that argument. You know, the clear cases that we have decided as a community are not edge cases. So we'll go with CSAM, because I have seen this in my own practice where a site has been compromised, and then you start seeing CSAM in your site. And as a -- And again, I'm going to go more to the small businesses than the large businesses, because large businesses have, again, different decision patterns in how they'll respond to something. And the small business may be -- it may be in the interest of the small business to say, "We don't need CSAM associated with this domain name. We market this domain name. We have SCO on this domain name." This has to stop now. We have to pause it. We have to reset. That's my point, that you cannot have a one-size-fit-all on compromised domain names. I think the best we could do -- and I'm sorry, I'm sorry speaking too quickly. I'm going to slow it right down. I apologize. That is U.S. East Coast just bred into me. I apologize.

What I'm going to say more slowly is there are nuanced decision trees that need to be made with compromised domains. We can't make assumptions about what a business owner may or may not want, and to your point, Alan, unless those cases are extreme.

But we shouldn't discount them, and we certainly shouldn't discount -- and I appreciate what Tucows does in that it's very communicative with its resellers letting them know what's going on.

I think a big part of the balancing problem is we have time consideration depending on how much harm is being done and where and then we have the consideration of the registrant who may very well decide it is in their best interest to have this domain suspended in order to reset it so that their own marketing, their own publicity around their domain is not, I'm going to use the word "diluted," which is a trademark term. But in this sense I mean that an otherwise good, respectable business entity can be ruined by the kinds of compromises that we see happen in our practice.

And I hope that's slow enough.

GRAEME BUNTON: Thank you, Lori, for that.

I have got Chris and then Alan. And then I have a question I would like to pose to the panel.

Chris.

CHRIS LEWIS-EVANS: Yeah, thanks. So I think for me, on our malicious domain, the registry and the Registrar Stakeholder Groups have done really good work to put our advice on the sort of evidentiary standards that they require for abuse notification.

And this work really hasn't been done for compromised domains. And I think here to allow for registries and registrars to act, there needs to be some evidentiary standards that need more explaining for the compromised domains.

So, you know, as an abuse reporter, are we able to say that we've contacted the hosting company, we've contacted the registrar, and there's been no action taken? However long that is, is up for discussion and again comes down to the harms being caused.

You know, focusing on our narrow scope that we take malware, you know, there is a large portion of ransomware going around at the moment. And a single infection of ransomware could be the end of business for someone and putting lots of people out of jobs.

So being able to articulate that in proper evidentiary standards, we recently had a large operation against Amotech which was causing massive harm to businesses and to individuals. Being able to articulate that and say we contacted the registrant, contacted the host provider, nothing's happened, we're now coming to you, the registrar or the registry, for a suspension because of this, really allows that decision process to come through.

And we also understand that then might have to turn into a discussion, you know? This is a large multinational company that you're asking for us to suspend. The impact of that is going to be massive. Actually, this phishing campaign is only targeting a couple of people. It has got a very small niche market, and the impact will be massive. Please, can we give the registrant another 48 hours to respond, and we will heighten our request.

So I think there's a large conversation piece here with a compromised domain. It really reduces the harm that's being caused by the compromised domain both on that primary factor and that collateral damage.

And, yes, it does require more evidence from reporters, but it also requires a bit more engagement from those in the chain taking action against that.

GRAEME BUNTON: Thank you, Chris. Lots in there.

In this particular -- and so I think you're totally right and that there's so much work to be done in this space. And it also very quickly bleeds past the responsibilities or expertise of registries and registrars into places like hosting companies. And that's a question I want to come back to.

But, first, Alan, please.

ALAN WOODS: Actually, after Chris, it makes my intervention really shorter because now I can just say one of the things -- I surprised myself and Chris. We're both on this, is SAC 115. And what we're talking about interoperability, making sure that the specific operator is engaged at the correct time, to ensure timely response, because our goal here is the time-to-life of that particular domain. The smaller we can get that, the less harm that can occur.

And by going to the wrong provider at the wrong time, it kind of enlarges that. And we need to try and avoid that. So shout-out to obviously the DNS Abuse Institute and future efforts there as well. Looking forward to that. Looking forward to being able to tie my back-end into your back-end and do all that.

But one thing I do want to point out that -- and I missed it as well -- is that some registrants, of course, are registrants of large platforms who they themselves have very sophisticated abuse monitoring procedures as well.

I think, Chris, you mentioned that a little bit as well.

We're not going to take down -- just to use the obvious one, aspersions being cast -- facebook.com. I'll say it. I'm not .COM. We are not taking down facebook.com because of abuse on Facebook. That would be ridiculous. So I think we need to be very balanced in that approach as well.

GRAEME BUNTON: Thanks, Alan.

So there's been a question that's come up a couple times, and I'm going to pose this to Alan and Reg again because I think this requires a little bit of elaboration for the community, which is:

What sort of relationship do you as a registry or registrar have with hosting companies? Because my sense is that there are many who think that that relationship is quite close. And I don't want to speak for you, but I think there's many circumstances where there isn't.

And so, you know, as we're talking about the complexity of this -- and Chris is talking about escalation paths between registrant, hosting company, registrar, registry, are those relationships clear? Are there standards for that process? Is this just a place where we should be doing some more work? Reg.

REG LEVY: Thanks, Graeme.

And, again, my answer is going to be predicated on the fact that Tucows is primarily a wholesale registrar.

So the answer may be different for other registrars. We don't have any hosting services, which is why we have Exact Hosting. And we have approximately 500 websites on that subsidiary of ours. So I essentially say that we have no hosting.

[Laughter]

Our resellers typically are also hosting companies. So it is frequently the case that we can reach out to a reseller and say there is a compromise on this site, and they can take care of without even involving the registrant. So in the event that our reseller is a hosting company, we have that relationship, and it is a close one.

However, that is not always the case. And there are many hosting companies. I mean, hosting is a service that requires a domain name. But it is completely separate from the service of providing domain name registration.

So in the event that our reseller is not the hosting company, it gets messy. And we need to use a DiG tool just like any other Internet user to figure out who the hosting company is and then rely on whatever information we can find only based on that DiG to contact the hosting company.

GRAEME BUNTON: Sorry, I lost my mouse cursor there.

Thank you, Reg. I will let Alan -- before I editorialize. Alan.

ALAN WOODS: Thank you. Apologies again to our interpreters. I know Irish people have fast speech.

From a registry point of view, it is much more difficult. We don't really have that connection with the hosting provider. We would probably expect that to be done before it comes to us. And also we would probably ask our registrar friends to see whether or not they could connect, if they are also the hosting provider.

What I will say, however, is that there is -- my dog is also snoring a lot.

What I will say is that we do engage with them, however, as much as we can outside. So obviously, there's several different conversations. One of those conversations is within the ICANN context, but there is one like the Internet & Jurisdiction, which we are also members of as well, because in that discussion, there are hosting providers who are also at that table because they're not within the ICANN context normally. We can have those conversations that build more bridges with the hosting providers at things like the I&J.

I think what's important for us, however, is to bring those learnings back to the ICANN community to inform what we are doing. I think that's something that we are doing more of. That's what happened with things like, you know, the Internet and Jurisdiction definition came into the contracted parties, and we worked with that.

So we're working with hosting providers as much as possible and bringing that back to the ICANN community is an important take-away for us as registries as well.

GRAEME BUNTON: Thanks, Alan. My sense is so many registrars are hosts but not all of them are. But determining the hosting or having a relationship with the host, given how many there are around the world, is not necessarily very common. And so that's a real impediment that we need to begin to think about, which is how do we improve those reporting process is to hosts. We're going to have far more tools at their disposal and be a more appropriate place for action. And then how do we define escalation paths for registries and registrars where there is an unresponsive registrant or an unresponsive host? And now we need to reassess an abuse complaint that has come on.

I have got a hand from Lori, and we have got about 18 minutes left. And so I think we are then going to move on to try and address some of the Q&A directly.

Lori, please, go ahead.

LORI SCHULMAN: Thank you. I wanted to make some points, again, following along in the chat as best I could.

But I do think in terms of reporting, interoperability, creating safe spaces, I do think it's worthy to insert into this discussion how much investment is reasonable and could be or should be expected into this space in terms of either mitigating technologies. And that could be humans.

To Reg's point, AI casts a very broad net. And EURid has been very open about saying and doing it's AI work. It still needs an army of humans to check the AI. Nobody takes the results for granted. They are expected. We understand for a safer Internet -- at least my constituency understands that a safer Internet may, in fact, require a lot more investment and as such, it may actually require higher prices. And do we want to talk about that? I know this is a concern of particularly civil society where it's important that the price of domain names stay relatively inexpensive and acceptable for anyone who needs a domain and wants to start a legitimate site.

But at the same time, we know that there's certain registries that -- and registrars that are investing more. Is that investment paying off? And I think that is an important question that the community could be asking itself.

GRAEME BUNTON: Thanks, Lori. In fact, I think that's probably a good segue for the next set of questions related to some of what's going on in the Q&A pod and in the chat, which is: What is the role for this community? Is there a role for ICANN in trying to address some of these challenges that we've discovered here today. It could be best practices for malicious registrations. What is the scope of ICANN's purview on compromised websites? What can we do there? What are the places in the community? I -- have lots I would like to talk about DNSAI, but I would like to leave that for a minute.

I'm curious if anyone on our panel has thoughts about where they would like to see this community go. I haven't heard from Rod as well. So I'm curious if you have got thoughts. Reg, you put up your hand.

Please go ahead.

REG LEVY: Thanks. I definitely think that the ICANN compliance team has a remit within the contracts to enforce the contracts where action is not taken against DNS abuse. That is to say against content that is within ICANN's remit that does not include, for example, illegal content, or CSAM as was discussed earlier, but for DNS abuse. And that ICANN should be taking better advantage of actually using those clauses and enforcing them.

GRAEME BUNTON: Thanks, Reg. I have got Rod and then Chris.

ROD RASMUSSEN: So I answered this question. There was a similar question in the Q&A pod which I answered. I think it was from Fab.

So SSAC actually spoke to this in SAC115 in that there is a role for ICANN org and the broader ICANN community, including contracted parties, everybody who shows up at these things to take part in the conversation. But it's a broader topic space and that we're focused on, you know, quote, unquote, DNS abuse. That's a subset of all abuse on the Internet. And some of the challenges we have we've talked to today around what are the appropriate service providers to be involved in mitigating whatever the harms are, whether it's compromise versus malicious is a great conversation in that space, but there's also -- think about evidentiary standards, how long -- or expectations about acknowledgment and mitigation of reported abuse, what kinds of things can be actions. There's a whole bunch of topics that are a broader thing.

We're seeing movement towards that. We have, you know, great initiatives like the domain -- domain abuse -- DNS Abuse Institute and some of the work that I&J is doing, Internet & Jurisdiction is doing. There's a whole lot of effort to create some sort of standards, best practices, et cetera, but we're not there yet, and we have a long way to go in trying to create a -- an ecosystem where people can have expectations both around process, proportionality, and the other items you need to create a better response and preventative system for all nature of abuse.

So it's important that we, I think, as the ICANN community be largely on the same page but then engage with the broader Internet community as well on how to deal with these issues, because if we all try and solve our own unique part of the puzzle, we're going to end up with a lot of different systems out there. But there is a lot of pointed work going on, and ICANN may have a role in being a place to foster conversations since they have resources and reach where a lot of other efforts, whether it's in hosting or email, et cetera, don't necessarily have that.

So I highly encourage folks to check out SAC115 and participate in some of the conversations that are going on to be proactive and think -- think globally about how to approach these problems and create those kinds of frameworks we need in order to, again, set expectations and follow through on them.

Thanks.

GRAEME BUNTON: Thanks, Rod.

I think I've seen someone put a link to the SAC115 in the chat already, but I'll make sure someone does that again.

Chris and then Alan.

CHRIS LEWIS-EVANS: Yeah, thanks. Chris Lewis-Evans for the record again.

So I agree with everything that Rod said there, but the I think I initially was going to say was to agree with Rich around the compliance. But I think, you know, we've (indiscernible) massively on how to deal with DNS abuse, and this conversation is part of that. I don't think we have a proper process for dealing with compromised domains which is a good part of the reason for this discussion. And having some sort of minimum expectations and having that documented so this allows clients to measure responses I think would be very helpful. And then also to have those standards rolled out to all the registrars and registries so they understand what is required of them I think is really key. And educational material for them as well. You know, there are a lot of different types of registrars and registries as has been said and understanding the best practice in how to deal with them I think is really key on top of this.

So there's a lot there that ICANN can do to spread that across the whole sort of ICANN landscape. And then just to build on Rod's -- Rod and do some outreach into other areas. You know, hosting companies, service providers we mentioned a lot during this, and they're really a big, key part to dealing with the DNS piece of it.

Thank you.

GRAEME BUNTON: Thanks, Chris.

Alan briefly, and then we'll try to answer a couple Q&A pod questions before we do a brief summary.

ALAN WOODS: Perfect. Very briefly, actually. It so happens that, you know, as part of this entire process, and even before we kind of suggest this plenary, the Registry Stakeholder Group DNS Abuse Subgroup has actually started a process to try and put out a paper, at least, to begin that. And Graeme is obviously involved in that. He's the unofficial leader at this particular moment in time of that, and I believe we're inviting people from SSAC as well. So we intend to bring in Rod and I believe Jeff Bedser so we can have this robust conversation about malicious versus compromised; to do exactly what Chris has just said there, is trying to lay the groundwork.

This was always an opening conversation. This was us acknowledging there is definitely an issue, we know, with DNS abuse and we're working on that, but we also need to work on those nuances and differences within there to help us effectively deal with it.

So I suppose watch this space, and hopefully we'll have something out soon on that as a very good and strong solid beginning.

GRAEME BUNTON: Thanks, Alan. So the paper we're working on is the TPH on this topic. I think we're really aiming to have that out between now and the meeting in The Hague, ICANN74 in June. I think it will end up being much closer to June than now, but we'll make sure that's circulated with the community.

Fab asked a question a long time ago in the chat about sort of this topic. You know, what can the community do? Are there changes to the RAA that can be applied to capture some of these best practices or these issues we're talking about today? And potentially. I think, of course, that's on the table.

I received a letter as the DNS Abuse Institute yesterday as a reach out from the DNSO small team on DNS abuse issues asking sort of this question. And I've been stewing on this myself a little bit, and I think the role for the community here is to really take some very small bite-sized chunks out of these bigger problems. We've talked a lot about the complexity in this ecosystem, especially where abuse is in compromised domain names. And I think we want to start with probably some low-hanging fruit on the malicious side where the consequences are smaller for getting it wrong and there's less victims involved.

And so I think we can begin to think about what work belongs on the community's plate for malicious registrations. And I think that also fits cleaner within ICANN's bylaws quite a bit. And so that would be my suggestion, and hopefully that answers Fab's question in the chat.

Let's see if -- we have seven more minutes. I want to give people an opportunity to see if they have any closing thoughts or if there's a question in the Q&A pod that anybody wishes to answer directly.

Reg has put her hand up. Please go ahead, Reg.

REG LEVY: Thanks, Graeme. I wanted to highlight something that Ashley actually said in the chat; that the registrar stakeholder group is currently working on a tool where you can put a domain name in, and it will spit out information about who the hosting company is and how to contact them.

So please keep an eye out for that. We premiered it at our Registrar Stakeholder Group earlier this week but hope to have live links available for everyone soon.

GRAEME BUNTON: Thanks, Reg.

And I think that's important and something similar that the DNS Abuse Institute is working on as well is going to be a centralized abuse reporting tool. It's pretty clear that those processes of identifying all the components in the ecosystem, who they are, how to contact them, what the standards are for contacting them, that's all pretty messy right now. And we can collectively do a better job of, you know, wrapping our arms around that sort of process and contactability problem and make that a little bit easier. And I'm working on some stuff like that. This isn't the place to talk about it, but I will be sharing news on that front later.

So let's see if we can answer a question or two really quickly here. Again, apologies for not getting to the is surplus of questions that we've had.

Maybe while I'm reading this queue, I'll see from the panel if there's any summary thoughts you have on this question that -- or this topic that we've been discussing here today.

I see Lori. Please go ahead.

LORI SCHULMAN: Sure. Thank you.

The chat has gotten so fast I can't follow it anymore. But I want to say I think this panel is timely and much needed. And I want to thank the organizers for inviting IPC and myself particularly because it does highlight the tricky issues.

I don't think anyone is saying any of this is easy. This is not an easy subject. I don't think anyone is saying that we should take rash decisions. But I do think that, at least from where I sit on the side of the house, is that we shouldn't -- particularly with compromised domain names. The community is experiencing what trademark practitioners and law enforcement and cybersecurity specialists have known for years, that each instance of abuse can have its own fact pattern. Each instance of abuse may have a remedy that's better or worse. Each instance of abuse is going to be looked at within its own set of facts.

So with all that being said, it's clear to me that we, as a community, it's incumbent upon us to establish the norms. And this is where I think projects like I&J and papers like SSAC115 and the work you are doing, Graeme, really help to establish those norms.

But where the next piece of this is, and I think it is appropriate to say how do the norms that are established outside of the community, how do they work inside of ICANN as well? And that kind of put a very far-thinking suggestion in the chat which has been discussed in my constituency, is when we had the issue of cyber -- sorry, cybersquatting come up 20 years ago, there was no judicial system in place or decider in place to help with these cases. And we found enough commonalities that we developed a UDRP, which has worked very well for 20 years.

Is it time to think about a UDRP-type process for compromised domains? And I'm going to leave it there, but I think it is a question for futuring, when we're talking about solutions and norms, it is worth.

GRAEME BUNTON: Thank you, Lori. Appreciate that input. And also futuring as a word is amazing.

I want to briefly, before I turn to Alan, answer a question from Griffin in the chat, part of which I think we already touched on which is the work of the -- inside the community, but also as someone who runs an institute dedicated to this topic and knowing there are several others out there, to me it is -- it's very obvious that abuse issues spill out of the ICANN remit very quickly where we need to be engaged with hosts. And so we have several organizations like Global Cyber Alliance, I&J, DNS Abuse Institute, topDNS from eco working in similar ways, working on similar problems. And what we need to do is be able to figure out how to work together. Collectively, I think we support the multistakeholder model and ICANN and really see, you know, clear importance of that work.

But we also need to understand that there's a place for adjacent organizations that can stretch beyond those borders and engage with the broader community in a way that maybe is not appropriate for ICANN or, you know, the rules within the ICANN system.

And so being able to explain to outside interests, be it, you know, regulators around the world, and to ICANN within the community that there's a place for both work inside the community that generates contractual requirements and PDPs and things like that, but also for industry to develop best practices, to reach out to adjacent things like hosting, email providers, that sort of community, that all of this can collaborate and work together. And we just need to get better at doing that, because it's going to benefit everyone and solve a bunch of these problems.

We have one minute left. It looks like, Alan, you get perhaps the last word.

ALAN WOODS: I just want to echo what you're saying there, Graeme. And I think it's important that just -- I don't disagree with Lori. In fact, I agree very much with what she said. I think we just need to make sure that UDRP is very much a domain name -- was addressing very much a domain name issue. What we're addressing here is something literally what Graeme is saying there, that is much broader than just registries, registrars. We're talking about hosting providers, that side of the ICANN context. So something broader than just ICANN, but definitely with constituent parts of ICANN, doing our part and working with others to try to come to a conclusion on this one. And I think that's literally the way that we are pushing at the moment, is that we are understanding this an interoperable area. We need to have the support and the understanding of all. And we're getting better at it. I do think we are getting better at it. And we must continue to improve.

GRAEME BUNTON: Thank you, Alan.

We're at the top of the hour. Thank you so much to my panelists. I really appreciate you guys taking the time. Maciej, for your great presentation. Audience, you've been wonderful for respecting our focus today and providing far more input than we could get to. I apologize for that. We'll try and capture some of that and see if there's a way to incorporate that into another session or some other work.

And so with that, I think we can end our session here today. Again, thank you, everyone, for the great engagement. I found that really valuable.

[END OF TRANSCRIPT]