DUBLIN – Spec 11 Security Framework Drafting Team F2F Meeting
Wednesday, October 21, 2015 – 11:30 to 12:45 IST
ICANN54 | Dublin, Ireland

UNIDENTIFIED MALE: Good morning, everyone. This is just to inform you that we will be starting the Security Framework Drafting Team meeting in two minutes. So if you would like to please take your seats, and if you need to continue your conversation, we would ask you please try to find space outside the meeting room. We would appreciate. This meeting will start in two minutes. Thank you.

UNIDENTIFIED MALE: It's October 21st and 11:30 in the morning in Liffy B. This is the Spec 11 Security Framework Drafting Team Face-to-Face Meeting.

FABIEN BETREMIEUX: Can I ask the members of the Framework Drafting Team to come sit at the table? If you're in the room and involved in the Security Framework Drafting Team, please come sit at the table. This will facilitate interactions as we expect that this session will be mostly an interactive session. Thank you.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

Good morning, everyone. Welcome to our meeting of the Security Framework Drafting Team. Our agenda today is going to be the following. We're going to very quickly go through a bit of background on this Security Framework Drafting Team. We will then have a series of opening remarks from the co-chairs of the Framework Drafting Team. They will then on move to present the approach to the drafting of the security framework, and then we expect to spend substantial time discussing some of the selected topics by the co-chairs.

So very quickly, a bit of background. Let me make sure that this slide displays properly. The security framework and its drafting team really stems from a series of decisions at ICANN, and in particular, the first one was the Beijing GAC advice, which introduced security checks as one of the six safeguards applicable to all new gTLDs.

Those security checks have two components: an identification of threats component and a response to identified threats. This was followed by a resolution from the NGPC in June 2013 that included some of the security checks in the Specification 11(3)(b). And also result to solicit community participation to develop a framework for registry operators to respond to identified security risks.

You hear I have the full sentence of that resolution on the screen because it should be considered fully, and I believe the co-chairs will point to that.

Between August and December 2014, we've engaged with registries to prepare the primary consultation that ensued where we consulted with registries and some representatives from the GAC on some preliminary thoughts around this security framework for registry operators respond to identified security risks. Primary consultation, which brought substantial comments, and led us to discuss those at ICANN 53 in Buenos Aires where registries volunteered to lead this drafting effort.

The objective of the framework ultimately used to reduce the impact of new gTLD-related security threats on Internet users through timely industry self-regulation. And I think that's an important component of the framework is really an instrument for self-regulation by the industry. And we expect the framework to become a set of reference sets of non-binding standards for registries, registrars, and all interested parties.

We formed this drafting team in July of this year. We, to date, have 30 registry representatives, five registrar representatives, and ten representatives from the PSWG, which is the Public Safety Working Group of the GAC.

The objective of the team, the drafting team, from our perspective is to produce the substance of the framework that would be grounded in industry experience, accepted best practices, and consultation with the relevant communities. And practically, we hope that the framework can be built collaboratively in the spirit of mutual agreement.

With this background, I'd like to introduce the co-chairs of the drafting team and start with Jonathan Flaherty, who's the co-chair. Do you want to say a few words? I'll let you introduce yourself.

JONATHAN FLAHERTY: Thanks, Fabien. Good morning, everyone. So I'll just give five-minute opening remarks, and I'm going to censor. I'll try and sort of set the scene of where I'm coming out of and from a law enforcement background, and why I feel it's good for me to be here today as the co-chair, and why I was nominated to be the co-chair for the PSWG on this working group.

And then I'll try and inform you in the mentality of thinking around voluntary frameworks by starting with something that happened 18 years ago in UK law enforcement. A different framework, but principles around the framework of accessing communications data with UK Internet service providers allowed me to educate myself with engagement with industry.

And the UK Government brief at the time in developing a framework to access data from industry came with a brief to give industry what we they want and what they needed for a framework to work.

First and foremost, it was all about engagement with providers. I learned within that engagement that not every provider is the same, business models are different, and capabilities and costs are different.

The guiding principles behind this framework that were produced by ICANN – Fabien just touched on a couple of them – marry really nicely into the framework that I worked on some years ago in the UK. And every one of you are different around the table, and the PSWG wants some kind of consistent response.

So the framework for accessing communications data came about in 1998. And in the 18 years since then, it's not changed, and if you think about how technology is changed, we must have got something right there in terms of cooperation and collaboration.

It's not a framework that could be copied because it's bound by legal process, and this is an informal voluntary best practice. So as the co-chair, if I come to this table and registries get a win out of the framework that works for you, that's a win for the PSWG. And it should be a framework for future dialogue.

So the communications data framework that's been written, we don't even look at that anymore in the UK. It just allows you to pick up the phone and ask for help. So that's what I'm trying to bring to the table and it's your framework, you know your networks, you know your brand and customers. If we can do anything for you today, it's probably going to be sharing best practice with you and maybe making sure we're as competent and consistent as we can be when we request for you to respond. Or maybe when we talk about our experiences of what you're doing already in terms of your diversification into cybersecurity products. You're not just a registry.

You might be talking about long-term mitigation, as well. So I'm really, really happy that I got chosen for this one. Thank you.

THEO GEURTS: Good morning. My name is Theo Geurts. I'm with the registrar group. I'm actually very glad what Jon just was saying here, the collaborative aspect that we are going to do in this frame working group here. So I'm pretty happy about that.

I'm also happy to notice that my registry colleagues, friends are on the table. I'm also very happy that contractual compliance is at the table here. Thank you. And this is… Let me touch on something here. Yesterday, we registrars talked about abuse extensively. And if you were at a previous session, you noticed a lot of speakers from very different entities, and they're also dealing with abuse.

So we, as registrars, we came up with a… I will call it a living document for now. It is not finished yet and it will deal about abuse in many, many aspects as we registrars are dealing with abuse day in and day out. And the variety of abuse is enormous.

So that document that we will present I think about in three weeks' time is going to be a starting point, a starting point to set up a broader discussion than just what is going to be discussed in this framing working group here.

How and what will be in this specification in this framework that is not known yet? But from what I just heard from Jon, the collaborative approach here, I'll welcome it. So if we can use part of other registrars are proposing, if we can agree on it, that will be great.

On the parts where we not can agree on, it will be really great to have the feedback from the Public Safety Working Group on why it is not acceptable. Because we need to tap into that knowledge of the law enforcement agencies. So with that, I'm going to throw it to the other co-chair.

YASMIN OMER: Thanks, Theo. Good morning, everyone. My name is Yasmin Omer. I'm from Neustar, and I am the registry co-chair. So I guess from a registry operator perspective, given that this is a framework on how registry operators are to respond to security threats, first and foremost, what we'd like to do is provide an understanding to the wider community as you would have seen in this week, there are quite a few sessions on abuse in the domain name space, and it seems to be a pretty hot topic at the moment.

So from a registry perspective, what we're hoping to do is provide an understanding to the wider community of what the registry operator's role is in the DNS ecosystem. And hopefully once that's done, we'll be able to also provide everyone with an understanding of the types of responses that are applicable and appropriate to registry operators when it comes to security threats.

It's also hoped that this framework and we'll go into detail later on in the session about what the scope of this framework is and what exactly we're trying to achieve. But it's hope that this framework will also serve as a tool for new gTLD registry operators and assist them in determining how – provide them with some options as to how they can respond to security threats. As you're all aware, there are a number of new industry participants, so it'll be pretty valuable to those new participants entering the space.

From the registry operator's perspective, what we bring to the table is the operational expertise. I shouldn't say just operational. We bring the real life day-to-day experience in the new gTLD space, as you're all aware, that there are differences in terms of the requirements and the volume of abuse that we've seen in the new gTLD space as opposed to the legacy space, and we bring that day-to-day specific knowledge of the technical solutions of the technical business, operational, and legal parameters that are applicable to a registry operator in this space.

And I guess finally – well, finally for now – what we're hoping out of this effort is that this serves as the start of the engagement with law enforcement. This is with – and I guess that the drafting team's been... I don't know, I think it's been a couple of months and I'm happy to say that engagement with the law enforcement guys and girls, and specifically the PSWG has been excellent. It's apparent that we do share the same objective and that's mitigating abusive behavior in new gTLDs. And yeah, the engagement has been great and we look forward to working, or continuing to work with law enforcement in this space.

So this document will hopefully start the conversation. It's definitely not the end, and will serve to provide greater awareness in this space and provide as a tool for registries to further engage law enforcement.

So in this drafting effort, we do have – we've come up with a number of collaboration principles and these are effectively – they do describe the roles of each of the three community segments. So we have registry operators, registrars, law enforcement, and members of the GAC participating in this drafting effort.

So first and foremost, the registries and registrars will be leading the drafting effort when it comes to developing the framework. The PSWG will be providing very valuable input on the framework based on their expertise, which we haven't necessarily had an opportunity to tap into in a formalized way before, so we're really looking forward to getting that valuable input and plugging it into this effort.

An environment of openness and collaboration, I know that sounds pretty cliché, but it has been the case, and it will continue to guide our drafting effort here.

One of the principles is flexibility in favor of specificity. So what that means, effectively, is that less is more. There are upwards of 1,000 new gTLDs. We all have varying business models. There are a number of different types of TLDs and this framework needs to ensure that it's universal such it's applicable to all the TLDs.

There's just no sense in having section of the framework that aren't applicable to a particular class of business model. So that's – it's a principle, it's a challenge, and Jon, I'm sure you can speak to this, but it is a challenge in putting together the framework, but that will also provide a level of flexibility and ultimately will be more beneficial to the recipients in the long run.

Jon, did you want to speak to the last two points?

JONATHAN FLAHERTY: Yeah. Thanks, Yasmin. So when we're in the practice, we react and respond to crime, [inaudible] in terms of cybercrime investigator. We try and get savvy and we like to work smarter rather than harder when we spot repeat requests into us. We look at enablers of crime behind single strand abuse takedowns, and what we find is that investigators in this field is they're working in a very – we're working through layers of ownership and lots of overlap between registries, registrars, and hosting providers to disrupt organized crime, particularly in line with the framework security threats of malware, botnets, and phishing.

So I think we've got to be nonprescriptive in how this document is written. The more flexible it will be, it's going to… We shouldn't be nailing ourselves to the cross at all on anything as an outcome here. The documents got to be articulated in a way that we still define responsibilities for these layers of ownership.

Yasmin's idea of the life of a domain in terms of its use in the ecosystem might tell you where everybody fits in in this process and when civil and criminal investigation agencies should engage in registry and when they shouldn't.

I think what did go on this morning, it's everybody's responsibility in terms of managing abuse. But sometimes, everybody's responsibility equals nobody's responsibility. So that kind of flowchart and those kind of diagrams in such a framework, they educate my members and me. They deconflict and stop duplicating the effort.

So I'm in total agreement with the collaboration there. But we'll see how that pans out. I think as [Theo] said in three weeks, when we start commenting, then we'll really see what the framework's going to be like, but the common shared interest there, [will win it through for us], I think. Thanks.

YASMIN OMER: Okay. So just we thought we'd provide just a quick background and provide an overview of the scope and limitations of the drafting effort, just so everyone's on the same page with respect to our objectives here.

So quickly, why are we here? We are here to – and when I say we, I mean the Drafting Team – deliver on the commitment made by the NGPC to the GAC regarding ICANN's soliciting community participation in a taskforce or through a PDP to develop a framework for registry operators to respond to threats.

I guess that's what's brought us to this point. It's not the only reason we are here. As I said earlier, it is serving as an excellent platform to provide that level of awareness regarding the registry operators, as well, and further engage law enforcement and registrars. Thanks.

So we're not here to be clarifying an existing obligation in the registry agreement. The obligations in the registry agreement relate to conducting a technical analysis and they're in Specification 11(3)(b). Maintaining statistical reports and providing these to ICANN. There [is] an obligation in the new gTLD registry agreement that relates to how a registry operator is to respond to identified threats.

So, effectively, this framework will be a set of principles and ideas that may be used by registry operators in deciding how they should respond to security threats with the general objective of mitigating security threats. And I cannot emphasize that enough. We do all have mutual objectives here and that's to mitigate abusive – well, security threats in the new gTLD space.

So what this framework isn't, as I touched on before, is a legally binding obligation on registry operators. To that end, the framework won't contain any SLAs.

It's not a set of rules that specify how a registry operator should respond to security threats. I think it's important to emphasize that it's – well, the distinction between the should and the may, again, the point I should emphasize in this respect from a registry operator perspective is that we're incentivized to mitigate abusive behavior in our TLDs. We're incentivized from the reputation point of view and, in some cases, from a commercial point of view.

So whilst there is concern for this being thought of as this framework being thought as a legally binding obligation, that's definitely not to say that we have no interest in taking steps to mitigate abusive behavior.

And finally, it's not a document that creates a presumption of compliance with the registry agreement. The point I want to make there is it relates to the previous point that there isn't a provision in the registry agreement that relates to how a registry operator is to respond to security threats.

So the scope of the framework. As I mentioned before, we hope to provide clarity on the registry operator's role in the DNS ecosystem, and that's something we'll discuss today.

Our responses to security threats, notification procedures, appropriate consequences, a discussion regarding ensuring respect is provided [inaudible] privacy and confidentiality. And finally, some case studies, and these case studies will really emphasize that responses to security threats are always going to be a fact- and-circumstances analysis.

It is really difficult to map certain response types to certain threats because it's the TLD differs the business model [inaudible].There are so many different nuances that impact how a registry operator responses to what particular security threat, and the case studies will hopefully demonstrate this and we'll certainly be seeking input from the PSWG on these case studies because it's anticipated that a few of them will include engagement with law enforcement and demonstration that that engagement has been beneficial to mitigating the security threat in that particular case.

So finally, I just wanted to emphasize that this scope is pretty much taken out of the NGPT, the new gTLD Program

Committee's response to the GAC advice. So it's important to ensure that we limit the scope of this framework to ensure that we deliver on this commitment as quickly and as efficiently as possible.

This definitely does not mean that registry operators – as I mentioned earlier , itdoesn't mean that registry operators won't continue to engage law enforcement registrars and the wider community in this space in particular. This is the start of the engagement. However, given that the reason we're delivering on this framework is because of the NGPC response, it's important that we do limit it to that scope.

And finally, some guiding principles in the drafting effort. We did touch on these earlier. The framework needs to be universal or not at all. No mapping of responses to threat types, the registry operator's policies do govern, I guess, the type of abuse, and you would have seen... Well, I'm sure you're all well aware that registry operators in the new gTLD space are required to include certain provisions in the RAA with respect to this type of behavior.

And finally, it needs to be cognizant of the role of the registry operator in the DNS ecosystem. So to that end, Fabien, I think we'll be discussing a few of these items for the remainder of this session. But if anyone has any questions. Oh, timeline? Sure. Yeah.

Sure. So in terms of the timeline, we hope to have a draft by the end of November. This is all tentative at the moment, so our registry colleagues, do not fret. By the end of November and for review by PSWG and registrar colleagues. We're hoping to have a few iterations of review. I'm aware that it clashes over – well, it's over the Christmas period, so we have taken that into account. But, of course, if more time is required, this is, by all means, a tentative timeline.

The hope is that – the goal, I should say – is that we have a document out for public comment during the ICANN 55 meeting and we're hoping to have some discussion with the community at that point regarding the framework.

Should we stop for questions before – yeah. Does anyone have any questions before we start discussions on selected topics?

Okay. So just very quickly, we'll start off with the typical abusive domain name processing by registries. That's effectively meant to be a, I guess, the life cycle of an abusive domain name in terms of our operational processes.

So Sean, yeah, can you please provide that update?

SEAN BASERI: Okay. Great. So this section will be educational and intended to provide further context for the subsequent sections. For this reason, this section will be a bit more broad. We'll discuss the process of receiving information about security threats, and this can include different methods. One common one is abuse reports, a discussion of analysis, and, for example, it's the relevant to the TLD? Is it covered by the policies of the TLD?

And then we will go into a bit into identifying actions, which is a later section. So we leave that for that just for those folks to discuss. And then actions and governance. For example, maybe notification to appropriate parties and other options.

YASMIN OMER: Excellent. Thank you, Sean. So next up, Jordan, just to provide an overview of the registry operator's role in the DNS ecosystem. Again, the intention her is that by providing this overview, we can provide an understanding of what the appropriate and applicable responses are for registry operators in this space. Jordan.?

JORDAN BUCHANAN: Thanks, Yasmin, and a little feedback. Jordan Buchanan with Google, for the record. And I'd like to spend just a few minutes talking about first what the role of the registry operator is as well as touching on the role of other folks in the DNS ecosystem, and notably the registrar in the DNS ecosystem. And to open up at least initial discussion about what that implies about the registry operator's role in responding to these sorts of security threats.

The important thing I think to start off to keep in mind is to look at the role of domain intermediaries in general. That's both registries and registrars. I'm sure everyone in this room knows it, but in the gTLD space, at least, we have a two-tier distribution model for domain names so registrars handle the retail relationship with the registrant and actually have a relationship with the registrant.

Sometimes, there's a reseller involved, as well, but for simplicity's sake, we'll ignore that at the moment since we're focusing on the registry operator role.

The registry operator provides a canonical database that provides some technical information like DNS servers and DNSSEC information and so on, and in most of the gTLD registry operators, also provides the database of contact information in the form of WHOIS. That's not true in a few of the older legacy TLDs. That's it, though. That's all the registry operator does, and it provides a registration interface to the registrar to fill that canonical database.

Notably, neither the registry nor the registrar has any control of what happens after the domain name points at something. The domain name itself generally is a pointer to some IP address or mail servers or various other types of information, but the domain name is just sort of helping the generally the Web browser but potentially other software sort of find the right server.

And so generally speaking, the registry or the registrar don't operate that server, they don't have access to what content is provided on the domain, etc. And so the implication of that is that all of the domain name intermediaries, the registry and the registrar, don't have – we don't have very many tools available to us in terms of responding to abuse.

We basically have three options. Really, there's two. But we can turn a domain off, right? So that's one option available to us. If we see sufficiently egregious case of abuse and it's in the right scope for our abuse policies, then we do have the option to disable the domain. But that's a very blunt tool. If there's one page on a website that's problematic or a particular piece of content somewhere on a server, when we take down the domain, it takes down all of the content on that domain – not just that one piece of information. So that's a very blunt tool that generally is only applicable in pretty egregious cases of abuse.

Generally speaking, in a lot of cases, you're going to want to try to get to the server, whoever's operating that, in order to try to target the response to the specific problematic either information or content, whatever it is that we're trying to target. In the case of some of these security threats, sometimes the domain might be used as to coordinate a botnet or something like that, and in those cases, the domain might actually be central to the operation of the attack.

But in most cases, you're going to see that if there's a phishing site or something like that, there will be a page hosted somewhere. Particularly, one of the problems we deal with here is malware, and with malware, often it's the case that it's on a website that's been compromised and so the owner of the registrant of the domain name isn't even aware that the domain is being used in conjunction with the distribution of malware, and they're an innocent victim, as well. So somehow, what we want to do is help them fix their website.

And so taking down the whole domain is actually – A, it doesn't help us fix the problem, and B, it's a fairly punitive action for what may be otherwise legitimate websites. But that is an option we have. We can take the domain name.

In very rare cases, there's been discussion of, as opposed to taking the domain down, pointing it at another set of – at a different server. Most of us, I think, don't have abuse policies that allow for that, and there hasn't been significant discussion within the community about when that step is appropriate. So in general, I'm not going to discuss that option particularly much today.

And then the last option is you can get in touch with the registrant, and that's often the right thing. Or figure out who hosts the server. But anyway, open up channels of communication, essentially. Because, like I said, we don't have control over what's actually hosted behind the domain name. And once again, here the registry operator's role in particular is quite limited.

We don't know anything more about how to contact the registrant than anyone else in the world does. We have the access to the public WHOIS information, but the registrar is the one that has the relationship with the registrant. And so, for example, if the registrant is using a privacy service or proxy service, the registry operator doesn't know anything more about that registrant than anyone else looking at the WHOIS record does. Only the registrar or the privacy provider would have access to that information.

So the net result of this, stepping back now with all of this information, is what's the proper role of the registry operator in fighting DNS abuse? In most cases, I would argue that the right place to be engaging, if you're engaging with a domain intermediary at all – in most cases, the right place to start is wherever the server is that the content's actually hosted on, because then you can actually target the specific problem as opposed to this one tool of trying to deal with the domain name.

But if for some reason we've concluded that the DNS intermediary is the right layer to act on, usually going to the registrar is the right thing to do because the registrar has that relationship with the registrant, they can make more nuanced judgments about it, they can get in touch with the registrant potentially in ways that someone with – above and beyond what someone with access to the WHOIS data would be able to.

And so in most cases, I think, if you get in touch with the registry operator to try to deal with a security threat, we're going to either say go to the registrar, or we'll just refer your report to the registrar so that they can handle it. So in most cases, it's probably more efficient just to go straight to the registrar in those cases.

There are a small number of scenarios, I think, where it makes sense to look at the registry operator being a proper role in dealing with these threats. So the first would be where there's massive, like large numbers of domains involved across the TLD or TLDs, if it's a registry operator that runs a number of TLDs. And it may not be efficient if there's hundreds of domains involved or something like that in a coordinated event.

Then it may make sense to go to a registry operator, where you could target action on all of those domains at the same time, as opposed to having to go to each individual registrar. And one example of that may be of this sort of large-scale action, maybe. I think Yasmin or Fabien mentioned earlier, there's a separate process outside of the security framework whereby registry operators engage in scanning and monitoring our own TLDs in order to understand what security threats exist there.

So that would be one place where we would make ourselves aware of security threats, and then we would want to, once again, generally let the registrar know about the problems so that they can take action with that with regards to the registrant.

But there may be other cases where law enforcement is working a case where there's a large number of domains or some other industry group would be aware of a coordinated botnet issue or something like that where we want to loop in the registry operator.

So that's one instance where it may make sense to involve the registry operator, and the second instance would be – and I think this one's a little more controversial, but we probably need to have more conversation around it – is where we've tried to engage the registrar in order to take action on a domain name and that's failed, for some reason.

And so then it may be that we need to say, "Okay, well that usual channel isn't working, so do we get the registry operator involved in that case?" And some registry operators already take action on these cases, some don't, so that's an area where I think some more discussion probably makes sense.

But beyond those two cases, it's generally speaking the right place to go is the registrar. And like I said before – and in most cases, you shouldn't be talking to either one, you should probably be talking to the hosting company or wherever the actual security attack is originating as opposed to with the domain name. Thanks.

YASMIN OMER: Thanks, Jordan. Any questions? Theo.

THEO GEURTS: Not a question, actually – just a comment and just to piggyback on the person from Google here. I don't want to make it more complicated than it is, and you are correct, in any sense there. That you should always go to the registrar to go to find out who maybe the reseller is or the registrant or the sub-sub-sub-reseller because there are many entities below there that a registrar is not aware of immediately. So we, as a registrar working with resellers, we often do not know who the registrant is also. We do not process their payments, we do not deal with their hosting, but basically, going to the registrar is always the right move because we can help you further there.

And to clarify it also, you're again was being just said. If you're dealing with a content malware, go to the web server. Go there. Thank you.

YASMIN OMER: Thanks, Theo. Benedict?

[BENEDICT]: Hey, Jordan. Do you take any industry data feeds like Spamhaus, [inaudible]? And if so, do you act on them?

JORDAN BUCHANAN: So that's more of a topic for the sort of scanning and understanding what's going on within the TLD, which is a little bit outside the scope of our discussion today. But that general – that does fall into the case, I think, one of the cases where registry operators are taking action and are at the right place because it's a broad gTLD scale large event, large set of data that you're acting on as opposed to individual domains

I think, as I mentioned before, all registry operators are obliged under Spec 11(3)(b) to perform scanning of their TLD and different registry operators have elected to do that in different ways. Often, that does involve third party data feeds.

Google not only consumes data feeds; we provide malware scanning through our web crawl and make that data feed available, so our own scanning relies heavily on our internal data as opposed to external data feeds, but I know that – it is a common practice among registry operators to consume third party data feeds.

[BENEDICT]: I'm sorry. I overstated. Not all registry operators are obliged to do that. Registry operators operating under the new form that have Spec 11 attached to it are obligated for the scanning.

MAXIM ALZOBA: Actually, we should be aware of situations where we cannot blindly trust all sources. Like, for example, with the Spamhaus, we sometimes have to contact them to eradicate our networks out of their lists because someone complained and it was blunt complaint.

Sometimes, we see flushes like someone is not happy with the competitors, I'd say – they flush information to all sources that their malware, they provide slavery, etc. And these things is not even in like two hours, but if we do accept, like cancel, the domain on full order to these sources, it's not possible. We have contractual obligations before third parties and, thus, this source of information should be manually processed or at least not processed blindly on full order.

YASMIN OMER: Thanks, Maxim. Guys, I just want to –the scope of the framework is limited to how we respond to security threats. Security threats may be identified through a number of mechanisms, but in terms of where this framework lies, it's after the fact. So it's what happens after we've identified the threats. So James, thanks.

JOE WALDRON: Thanks. Joe Waldron from VeriSign. So thanks for the clarification, but that wasn't what I was going for. I did want to clarify one other point. I think you used the word scanning, which is not the term in the agreement. It's pruritic technical analysis. While scanning may be a part of that, I think scanning has a completely different connotation in the original intent of technical analysis. I just wanted to clarify that.

And I would like to just tack on an additional point that I think Jordan explained very clearly in terms of what the limitations are of the registry. And I know one of the concerns that we've had as we've taken actions on names for many years, is really based on that limitation that Jordan described, and it really is the unintended consequences of taking a domain name down where we don't necessarily know some of those services. So I'll use a hypothetical example.

If somebody wants to take a domain name down because of spam, and there's one e-mail address on let's just say a worldwide e-mail network, I don't want to take the entire domain name down because of one e-mail address.

We've had similar situations like that where a domain name that has child name servers that's hosting, in some cases, millions of domain names, and you take down that parent domain name, you're impacting potentially millions of other domain names unintentionally.

So that is a limitation not only of registries but often of registrars, where you don't know the full impact of some of the actions that you take. And that's why think we need to be very cautious in the approach that we define. Thanks.

YASMIN OMER: Sheri?

SHERI FALCO: Hi. Sheri Falco, ICM Registry. And sort of on that point, one thing we noticed in the .xxx space is a lot of the… We use Google's malware sort of scanning services well just to kind of get a sense of what's going on there. And one of the things we noticed early one was that a lot of the providers in our space use advertising programs and it's often the advertisers, so third parties not even associated with the domain name registrant. And they're the ones that are sort of showing up and triggering the Google notification. And so it's an important tool for us to then reach out to our registrants and communicate so that they're aware.

Because often, they might not even be aware that that's occurring. And, obviously, they have an interest in fixing that for their own purposes, as well, but a lot of times the malware is not even being triggered by the actual registrant, so taking the site down would clearly not be an appropriate response. But figuring out a communication structure and an escalation structure around that has been something we've been using.

YASMIN OMER: Okay. So we'll move on to the next topic, and that's responses to security threats. Alan Woods from Rightside.

ALAN WOODS: So this one is where we're obviously getting into a bit more of the direct and a lot of issues can come up, and when we're talking about specific responses. So we're trying to avoid the concept of going to specific responses.

So because we have a myriad and so many registries, so many different policies and procedures, you need to look at it from a much higher level and much more universal principle. And instead of looking at a uniformity of a response, one should really focus on getting a uniformity of the goal. And that's what we need to focus is that we all might use different ways of getting to that goal, but as long as we end up in relatively the same place, that's the important thing.

So with that in mind, putting together high level principles, and because, as you can tell, I'm from Ireland and Dublin, therefore, I'm painfully European when it comes to things like principles and proportionality. And I like to think along the lines of that principle of proportionality, and that is that the minimum actions causing the maximum effect.

Therefore, you don't want to have – again, taking into account over-acting can cause unintended consequences, as already been pointed out. So we're looking at very simple things, in the spirit of collaboration, as well. So responses should – things such as we need to focus on that we know that each registry, each registrar have that single point of contact. It is a requirement for registries on abuse, which is part of the topic, that we have that single point of contact.

So as a high over-arching principle, we should know who we need to talk to so we can quickly and effectively talk to that person as is necessarily. I mean, that brings into then, obviously, who is the appropriate party? Any one of the people in the registration link or in the abuse timeframe or timeline or the security threat timeframe.

A report can be received by anybody, but we need to know at a high level who's the appropriate party to actually deal with that. And again, Jordan was talking about that in many instances, because of the proximity to the registrant themselves, that would be, of course, the registrar. But that does not say that there is a part to play by other people in it because we all have the reciprocal responsibility of then aiding the registrar, if need be, in an instant.

So again, with that proportionality, that we all need to be prepared, in which we are, that when there is a part to be played, when it is appropriate for that party to play that part, that we will play that part.

Another very high level principle, as well, is that it must be timely response. I mean, security threats, obviously, are linked to time-sensitive issues. Therefore, again, all the people, all the parties in the chain of the mitigation of this need to be aware that the faster and the more efficiently something is done, the less impact that a security threat could potentially have or the less effect that it could snowball into something a lot larger.

So again, very high level, very common sense, but again, if you're looking to get to that specific goal, we're not looking at individual policies and principles for a particular registry to do, but what you may do and may take into consideration when you are forming those principles yourself within.

Another one, as well, is – and I put this – is justification and transparency. Again, if actions are being taken, there needs to be a justification for that action, and therefore, at a high-level principle, it's a matter of good – in my mind, it's a matter of good business, it's a good matter of abuse management, it's a good matter of compliance just to have proper records, proper reasoning that if ever you were called to task that you can turn around and say, "Well, these are the reasons, these are why we took such an action."

And again, that applies to registries, registrars, resellers along the entire line. So again, high-level, all aiming at that same goal.

So to move on to the notification aspect of it, then again, painfully connected to proportionality, much the same sort of concept. We need sufficiently identified contacts and we need to communicate appropriate with the appropriate parties.

So if you are responding to a security threat, that notification must be properly directed. There's no point in having the machinegun effect of let's send 17 e-mails to 17 people, hoping that one person will take responsibility for it. We need to know who and to whom it's going.

Again, simplicity in the notification is that they should be detailed and they should be clear. You can't just say, "Spam..." Well, not spam. Okay, "Phishing or malware attack at this domain. Thank you." You need to have the specificity. You need to have the detail.

So again, on a high-level principle, there should always be that clarity so that when the appropriate party is reviewing, that they have an action that they can take, or at least they have the tools to be able to make a decision themselves or to pass it on as appropriate.

Again, in notifications, there's nothing worse than notifying and finding there's a black hole. You need responsiveness of the parties, and again, this is what I think this is about. It's collaboration. We know exactly that if I am going to be communicating with, say, other registry [inaudible] communicate with the registrar, that that's not going to fall by the wayside, that that notification will be taken seriously, considering I'm giving the right and proper notification of not – I'm providing the information properly. So again, responsiveness is very important.

Another area, which is, again, a very high-level principle but, again, equally important is that we are not necessarily – because these are high-level principles, we're not limited to saying in X instance, you must go to the registrar. You can't limit – each party may have to jump specifically depending on severity, depending on source to other parties.

So be that in very limited circumstances, maybe to the registrant directly, but maybe to ICANN, maybe to the authorities. But again, it is – we can't say in this framework in what instances should you consider that. Again, it comes very much down to the individual registry, registrar, reseller to make that call at that time based on their policies, based on their principles. But again, ultimately, with the common goal, and that is the mitigation of that security threat.

So that is where, I suppose, our thoughts are at the moment. High-level principles are all aimed at that one singular goal.

YASMIN OMER: Thank you, Alan. Any questions?

UNIDENTIFIED MALE: We have question from remote participation. It's Nick Shorey from... You're here?

NICK SHOREY: I just [inaudible] sorry. I wanted to know your name. I didn't catch your name.

ALAN WOODS: Alan Woods, Rightside Registry.

UNIDENTIFIED MALE: So if we can repeat the names and participants are asking for names for us, so it will be good to repeat your name beforehand. Thank you. No. Just the speakers.

YASMIN OMER: It's just a reminder to please announce yourself. Thank you.

KRISTA PAPAC: Krista Papac, ICANN staff. Taking my own advice.

YASMIN OMER: Okay. Great. So that wraps up the registry discussion of selected topics. We shall move on to the PSWG's topics.

JONATHAN FLAHERTY: That's really, really good, Alan, to – you've just stolen all of my thunder, by the way, to the first few points. So we're the requesters into you, and we have a lot of best practice that we're aiming to put into this framework to meet lots of human rights considerations to make sure if a request goes in, it's necessary, it's proportional. We carry out things like sanity checks on the domains and we do due diligence around them to make sure there's minimal risk in taking Google.com down and we've had some stories like that. I won't mention the agency, I think, in the US in the past that did something around that.

So we have a policy behind that. I'm referring to the National Crime Agency policy. It will overlap, I'm sure, into registrar fields and speaking to Michele in the week, I'm going to e-mail him a lot of that documentation.

We work on a single point of contact approach to abuse, so the application, the notification, and response is very consistent. We can even put forward as a PSWG, we can access a referral function, if you're not happy from our side on request or anything like that.

You want a consistent approach there. So we're wired into this process. We're geared up to react to crime and to have some measures in place that hopefully make your lives easier in terms of responding to maybe what might be day-to-day security threats.

We would see that usually going to registrars and registrants, as well, of course, if we're looking at compromised infrastructure and maybe we need to contact a victim of crime there. So there's lots of areas there, these layers of ownership mean that there might be more than one party that we speak to.

Benedict's mentioned timely security threat feeds. The submission I've done so far for this working group is an options paper on things that we think constitute potentially the way a registry already is, maybe, or could look into in terms of responding appropriately or periodically to security threats. One of them was timely industry self-regulation against the ICANN guiding principle and daily feeds from a variety of companies.

So off of that industries, you could outsource that. I'll probably have to namedrop shadow server because that was the example in the options paper I produced. And they're free, we shadow server, so they're tailored to TLDs and they're in the format that you, I think, would want. They're very, very timely, current, actionable intelligence reports from what I've seen.

It could lead on to then your wider picture, as I see, as a registry, in terms of the overall threat. You don't want necessarily to be bombarded with single-strand takedown requests that are poorly put forward, and you can see more of the space at the registry level.

So in terms of threat sharing, I think that framework for dialogue. My liaison in the last two or three years with registries has been ringing up to, actually, about a problem rather than expecting a response. I'll just give you an example. We look at target static IP infrastructure and some scenes of crime, which in the 21st Century, are often a computer server.

The seed of the investigation often starts with an IP address in a log file [inaudible] of a network intrusion, for example, or some malware hosting. And we'll break that out and we'll branch out on that. Normally, the IP leads us to a domain name, and we wider map that infrastructure then. We want to look at whose responsible for the IP block and what domain name traffic potentially has pointed to or transited the IP. And in one instance, we had – we stumbled across really through passive DNS of an IP range in question – what looks like DNS tunneling. What looks like, which is legitimate way of passing data and content over port 53 rather than port 80. But also can be used potentially criminally.

That [inaudible] the threat picture for us. We've seen that on the Home Depot, one of the biggest hacks of all time. I think there's 56 million credit and/or debit cards siphoned out using point of sale malware and DNS tunneling was the method of exfiltration for that.

That's what we'd like to share with you, and I'm going to take Jordan's point and I'm going to corroborate with what he said. It's not necessarily anything to do with a registry, but you might be interested in taking that on, and it might be a bit of a threat picture that you didn't know about.

So we kind of send you a report and expect a report from you when we least expect it. It's like an ISP seeing a spike in traffic or some trigger on a voluntary basis might mean that law enforcement can offer a little bit more on this. I think there's a limited number of agencies that can go in at that level, and a lot of them are on the PSWG. So working with those security threat feeds, marrying that to our own investigations might show you something that's new and coming that's around the corner in terms of reporting. So threat sharing has got to be a part in the framework.

Quarantine of the domains via third party registrar was another item in the options paper. Infrequent, bespoke, complex requests are what we sent to registries. In a no one-size-fits-all industry, we often ask questions of sinkholing from registries – specifically, again, I seem to be talking about cybercrime in full. I think cybercrime investigators are probably your biggest customer here.

It's a burden on your resources. I think the ability to refer that to – I shouldn't really say the name, but a third-party registrar who might, in the future, be able to do that, would be – it's kind of like law enforcement organizing this to tailor to your needs, and a solution, as well as the problems that [will go in] so that every registry has a chance here to refer. You may wish to do that yourself. So a kind of referral system for that would be good.

The majority, I think, of the responses to security threats are going to be probably botnet sinkholes. The prep that the simultaneous sinkhole [inaudible] takedown of that has never been more apparent for me in what we did in September in terms of the takedown of the Dridex malware and sinkholing via third parties, speaking to ISPs on a consensual basis to transfer IP addresses to ownership of a third party by consent to then sinkhole.

Got some really, really good results, meaning at a given time, we'll be working with you in a joint investigation, joint industry, outsourced where needed investigation. The high-end requests that go in, those are, in a framework, especially one that would go for public comment, not really to be itemized and be prescriptive too much. I would class them as a special service and not mention too much about them.

And should we have a special service request like that, then the theory and the framework is, yeah, a further meeting takes place and we take that from there. So we've got lots of ideas and we try and take that to you with options, as well, in terms of how you might want to manage that.

Long-term mitigation at scale. I'm still theming this on cybercrime. Mitigation of scale is a big concept amongst civil and criminal investigation agencies now. We do a lot of whack-a-mole in terms of botnet investigation. I think a lot of it is as fast as we can make it, and along with getting an arrest on the end of it, that's the ultimate deterrent.

Looking at botnets just as a long-term mitigation at scale example, the last four jobs across different malware strands we've worked on all show commonalities. And they all start with a spam campaign. Cybercriminals are very organized and they carry out domain name-based spoofing, their spam campaigns. And they used to carry them out against banks.

I can't remember the last time I got spam e-mail from a UK bank. They just don't come in the webmail box anymore. They used to be in the spam folder for a bit and they're not there anymore.

So eyeballs on for the criminal now from what we can see on these jobs is that they target retail. They're the next set of financial targets where the chances are someone's going to open an invoice from a retailer. And they specifically are targeting UK domains and long-term mitigation at scale for me is if you stop the spam, you stop the malware, and you haven't got a botnet in the first place, which is easier said than done.

So we're working on a DMARC campaign, it is more aligned not to the registry, but it's to an e-mail provider, but it's all about domains. Looking at Alexa traffic rankings, we've just taken… We've DMARC checked to see who's got e-mail authentication on a domain to make sure a domain [and the front heade] are matching.

Of the top 6,500 most visited UK sites, and there's about a 6% DMARC uptake across the sites, mainly there's a poor show in the retail sector. So we're going to launch a campaign there and at the UK end. We want to do that if it's successful globally. I know Google and PayPal are heavily behind DMARC, so Google are not just a registry here, they're a company of many arms. We want to get into INTERPOL involved in a global campaign for that.

And it might be that DMARC’s not right for everybody, and I’d get you involved in that and I’d share that with you. So those kinds of views on the world, if you see anything that… A change in a spam campaign. If you can see that or in another subsidiary of your company can, I’d really like to do something about that in a non-pursue, like a protect kind of area of crime.

DNSSEC might be another campaign. And a note from Benedict that the SSAC have some educational resources that they’ve produced for the registrant in how to implement DNSSEC, a new secure-by-design, secure-by-default standards like DANE are on the market to provide TLS security on top of the overlay of DNSSEC.

We’re trying to up our game on the PSWG with the view that long-term mitigation of scale ultimately just allows you to carry on protecting your brand. If you do report to us, we intend to do something about that. We don’t want you to report us for the sake of it.

We’ve got all of these ideas and if you take one or two of them on, I think it puts us on another platform in terms of cooperation. We’re not using you as a service. We’re using you as kind of an ally in the constant arms race of cybercrime, so I'll stop there. Thanks.

NICK SHOREY: So thank you very much. I'm Nick Shorey. I'm part of the UK's Government Advisory Committee Team with a kind of a focus on the Public Safety Working Group from that angle. And I'd just like to pick up on a point that Jon made there about the wider piece and sort of implementation of DMARC.

So in our role, we're sort of focused on ICANN, but we're also involved in the Internet Governance Forum, and there's a best practice piece going on there at the moment. We're also looking at this sort of mitigation of unwanted communications and spam and that.

And so that's one of the sort of the key sort of opportunities. When they say best practice, they're not sort of, it's not enforcing anything, it's just sort of ideas. Again, that kind of melting pot of ideas. So if – and I'd sort of recommend that. I'm sure many of you are involved within the IGF already, but if you're not, definitely get involved and participate in that, so thank you.

UNIDENTIFIED FEMALE: Okay. Jordan. Just in time.

JORDAN BUCHANAN: Sorry. Hopefully this will be quick since we've got one minute left. Jordan Buchanan from Google for the record again. So I wanted to just go back to the quarantine option for a moment just to better understand are you actually talking about transfer to a third party registrar like in the ICANN sense of like a transfer of sponsorship?

UNIDENTIFIED MALE: I'm guessing I should probably answer that. Yeah, so to be clear, the quarantine known as the registrar of last resort is an ICANN-accredited registrar. We've got an IANA ID quite recently, which is pretty exciting to me and probably no one else in this room. And we're also accrediting with ccTLDs, as well, so the slightly, the nuance that's worth articulating here is that a lot of bad domains don't exist, they're not used by criminals until – they're only used in potenture, and they registered maybe a couple out of 1,000.

So in that case, we'll be either working with registries to add them to a block list, an internal block list, so do not register these domain names, and I know a lot of registries do that internally anyway. So we can be a mechanism for that.

But also to transfer, so force transfer, existing registered domain names under the bringings and due process to that. I suspect the ICANN policy we'll be using for that is the little-known ERSR, which is generally under court order within the US. Is that clear?

JORDAN BUCHANAN: Yeah. That was really helpful. I was just going to say, so this reminds me, I think it would be a be a, but possibly the ICANN meeting before that, there was a session on domain hijacking, as well. A little bit of a cybersecurity problem, but more focused on the registrar space once again.

And I think the conclusion from that meeting, and I think thinking about this option, as well, leads me to believe that there's probably some policy development work that the GNSO should do to enable more rapid transfers between registrars in certain situations. And so I think probably both the registries and registrars out to talk with their GNSO Councilors to think about initiating some policy work in that area.

FABIEN BETREMIEUX: And before we close the session, we have a question from a remote participant, [inaudible] from India. "Does the same mitigation techniques are completely applicable or enough to respond to security threats while dealing with IDNs? Does IDNs exposed to additional threats that might not be there with other ccTLDs, gTLDs?

Would anyone to address this question? Sure. To you.

THEO GEURTS: Okay. So I don’t have the numbers. So from what I’ve seen in the report, there is no correlation there or a higher uptake.

FABIEN BETREMIEUX: All right. So I think that completes this session. Thank you all for coming today and for your discussion, and we’re looking forward to more engagement in the development of the framework. Thank you very much.

[END OF TRANSCRIPTION]