JOHANNESBURG – ALAC and Regional Leaders Working Session Part 5
Wednesday, June 28, 2017 – 13:30 to 15:00 JNB
ICANN59 | Johannesburg, South Africa

UNIDENTIFIED MALE: This is the ICANN 59 ALAC and Regional Leaders Working Session Part 5 on the 28th of June, 2017 from 1:30 to 3:00 in Ballroom 4.

ALAN GREENBERG: All right, thank you. Welcome you to ALAC Session #5. We have three different speakers, 30 minutes each and we're on a pretty tight schedule so I'd like to start right away.

Our first speaker is Greg Aaron. We're going to talk about – excuse me, I'm out of breath – talk about domain name abuse. And Greg's one of the people on the forefront of watching ugly things that are going on, on the Internet, and hopefully trying to protect us. And I think it'll be an interesting talk. I'm not going to waste more time on introductions.

GREG AARON: Thank you, Alan. I'm a cyber security expert. Basically, one of my specialties is observing how cyber criminals use domain names. I work for a cyber security company called iThreat. I'm here in that capacity. But I'm also a member of the ICANN SSAC, and I'm a Senior Research Fellow at the Anti-Phishing Working Group.

And let's move on to our first slide, please.

One way to define domain abuse is some activity that requires or uses a domain name to perpetrate harmful activities. Now, within that definition, there's some wiggle room. But I tend to concentrate on some phenomena which are very well-known and are a core set of cyber crimes. And each of these uses domain names in various ways, but their domains are fundamental to carrying out these activities.

So at the center, I have spam, and the reason why I put that there is it enables a lot of the other kinds of problems. I define spam as, in a traditional way, which is bulk, unsolicited e-mail. Now, in some jurisdictions, that may not be strictly illegal or so forth, but I'm talking about activity which is illegal in most places or it uses patently illegal means.

85% of the e-mails sent in the world is considered spam, and the majority of that is sent from botnets, which are networks of compromised machines that have been hacked into and infected with malware. So sending spam in that method, that's actually a very bad thing. That is a criminal kind of activity to build a botnet or to use it.

Spam is also used to advertise all kinds of things. Some of it is fairly innocuous. What happens is these are… we're talking usually about domains that are advertised in the body of the

mail itself. This is the place that the spammer wants you to go to, so that's in some sort of a link in the mail itself.

And that consumes, actually, a ton of domain names every year. We're starting to track exactly what that means, but it seems to consume at least 8 to 10 million unique domain names every year. So out of the 300 million domains that may exist in the world's registries, there's a percentage of it which is involved in this activity.

Spam, for example, is used to advertise phishing. Phishing is where a criminal sets up a website that spoofs or imitates a site that you might trust, like PayPal or your bank. They want to get people to go there to insert their credit card numbers, their names and other sensitive information, which is then stolen and used to defraud people, drain their bank accounts, and that kind of thing. So that's how most phishing is advertised.

By the way, a lot of this stuff you don't see in your mailbox because there is a legion of people and services trying to keep this out of your mailbox. But a lot of it will slip through.

Spam is also used to advertise a lot of fraud, what we sometimes call 419 scams, like the Nigerian prince who wants to send you lots of money. And spam is one of the ways in which malware is spread. Again, somebody wants you to click on something and that's going to lead you to a site where you download something

whether you know it or not, and that's going to infect your computer. Some of that malware, of course, gets you hooked into a botnet and your machine is used for activities you may not know about. Botnets, among other things, send spam.

They're also used to launch DDoS attacks. These are Denial of Service attacks where tremendous amounts of traffic are sent to a destination in order to disrupt it and bring it down. And DDoS activities are rising and the amount of bandwidth that they consume is rising a lot. You may have heard about Internet devices like video cameras, or you know, security cameras and appliances being used for that kind of activity because they're infected with malware.

So this is the core set of activities I usually talk about because they're pervasive, they're fairly well understood, and they're basically criminal activities. And I think there's good consensus around that kind of thing pretty much everywhere in the world. Next slide, please.

This is an example of the kind of tracking that's being done in the security community. This is tracking I do with the Anti-Phishing Working Group. And this looks at phishing attacks and the use of malicious registrations of domain names.

So you see the statistics on the number of phishing attacks, in other words, these particular pages or sites that are used to

perpetrate phishing, and it's grown about fivefold since we began tracking it in 2008.

The number of domain names that phishers are also registering for themselves is also rising a lot. 2016 was the first time where we saw domain names being registered in such huge quantities. Phishers can break into websites. Website owners and hosting providers are also victims of these kinds of activities, and they can put phishing sites on somebody else's website. That's actually how most phishing was done, by breaking into somebody else's website.

But what we're seeing in 2016 was phishers are just going out and buying domain names. They always have done that and a lot of this activity is certainly located in certain places in the domain name industry. Certain registrars sell a lot more and some of the registration activity takes place in certain TLDs. And sometimes that changes over time. That has always been the case.

But a lot of this activity is concentrated, actually, at a few places. And so mitigation becomes an issue of some providers that are vulnerable in some way or another. Next slide, please.

So I wanted to run down kind of just some of the realities that not only security practitioners or responders have to deal with, but kind of how the ecosystem itself is organized.

The bottom line is that these kinds of activities are very pervasive, and criminals that are perpetrating them are often very professional. There are some amateurs out there, absolutely. But a lot of the people who are doing this are making millions of dollars a year. And the dollar amounts that are being stolen from individuals and companies has generally increased over the years.

A lot of company bank accounts, large and small, are targeted. And we're not talking about a few hundred dollars on people's credit cards anymore. In some cases, we're really talking about tens of thousands of dollars being more the norm.

The abuse does tend to concentrate in certain places at certain registries, certain registrars, certain hosting providers.

Why is that? Sometimes they're not paying attention and they may not have somebody on their staff who looks after these problems. There's a problem with prevention, which is knowing your customers and keeping out people you don't want to be there. And then there's the problem of once there is an issue and it's reported to you, what do you do about it?

There's the issue of whether you're proactive as a company to keep these kinds of activities under control or you respond to them when people tell you or do you not respond at all. And attention is one issue.

In the domain name industry, low price is absolutely an issue. One of the problems with the new TLDs is that domain names have been very cheap. There is a lot of competition in the industry and millions of domains have been sold in the new TLDs specifically for spamming. So that's an issue that attracts criminals because they don't want to pay as much money for things just like the rest of us.

Every once in a while in one of these industries hosting and so forth, there are complicit or even criminal actors. In the registrar space for example, there have been several registrar businesses that were owned by criminals for the purpose of perpetrating their activities.

Two of them were [S Domains] and ABSystems. [S Domains], the owner was convicted of wire fraud and some other charges in Estonia. The owner of ABSystems was using his registrar to spam out an operate domains that sold illegal pharmaceuticals. After he was finally arrested by the American Drug Enforcement Agency, he admitted that he had put out contract kills on at least two people.

These people actually exist, and my point here is they actually know the domain name system. Sometimes they're actually in it and they will not play by the rules of normalized society.

Mitigation is mainly done by private parties on the Internet, not law enforcement. Law enforcement is incredibly important. However, they don't always have the resources they need to pursue things and, of course, they have to pursue large cases and not smaller cases involving smaller amounts of money, for example. They need to consolidate cases that have affected tons of people.

And so they're overwhelmed and they work very hard, but the reality is we cannot rely on them to protect all of us because it's just not possible.

On the Internet, instead, relationships tend to be governed by contracts. If you want to use a service like Google or Skype or you have a terms of service for your mobile phone, and your cable service, and your ISP. These are the contracts that tell us what the terms of services are and they also give a party the right to shut off your service, for example, if you're violating those terms of service.

What tends to happen is people in the security industry or victims, they go and they talk to the service provider who's involved. So they'll call a registrar or a registry and say, "There's something bad happening here. Will you consider taking an action?" Or they will call the hosting provider and say, "There is a phishing site here. Can you take that page down?"

So there is cooperative enforcement, basically, and the companies have to decide when and if to enforce their own terms of service. People also choose their services based on terms of service in a competitive environment.

My personal feeling is, and that of many people in my industry, is that anybody who's operating an Internet resource has the responsibility to do so in a wise manner, which takes victims into account. The Internet is a network of networks. For a lot of its activities, there is no governing body or anybody who kind of oversees things, and so we do the best we can. The Internet is open and that's a wonderful thing, but cyber crime is one of the drawbacks of that decentralized system.

Next, please.

So the question does come up, what is ICANN's role in all of this? And this is some of my personal thinking.

The Bylaws do say certain general things about what falls within ICANN's remit, specifically to ensure the stable and secure operation of the Internet's unique identifier systems, and policies for which the uniform or coordinated resolution is reasonably necessary to facilitate openness and so forth.

The question then becomes, well, how far does that extend? What does that really mean?

One of the things that ICANN does do is it credits the registrars and the registries. It's saying, "You have the right to operate a registry and to sell gTLD domain names." That implies permission, but then ICANN can set some policies which are placed in the contracts through our community processes and so forth.

ICANN does things like access to WHOIS data and zone files. Those are two extremely important tools for security and protecting people. That data can reveal a lot of information about what's going on, on the infrastructure.

There are prohibitions in the contracts against using domains maliciously. Registrants have that in their contracts. Anti-abuse monitoring requirements for registrars and registries are in there.

Ultimately, ICANN's contracts are enforceable contracts. If the provisions weren't enforceable, they wouldn't be in there and they wouldn't be contracts.

One of my suggestions to people is you have to concentrate on the biggest problems. As I said, a lot of these malicious registrations are made in certain places. Understand why that's happening. Encourage the situation to get better in various ways.

So that brings me to the end of the slides and I think we're going to have some conversation.

ALAN GREENBERG: I hope so. Andrei.

ANDREI KOLESNIKOV: Thank you. Greg, I think it's very important to say that the domain name abuse is a part of the much larger international cybercrime economy, which accounts for billions of dollars. And it's not just the guys who is abusing domains or do bad things. That's a system which is a kind of dark economy with billions of dollars.

And I really believe that having online tools which help registrars to get the actual data on their domain name use goes beyond gTLDs. It is also very important for the country code domain names.

And of course, you know that .au, I was, when I was in charge many years ago, launched this nice online tool which people can really check their domain names and registrars can do the bulk check if they receive the reports.

But to do this, it requires some funding, of course, some resources like big registries in the countries or gTLDs can afford creation of such a tools.

And we actually did it on a non-commercial basis, funding these tools out from the revenue flows from the domain name registration as a registry of the national registry.

How it is organized, I mean, other services for the registrars and registries who can use, like these kind of online tools, a bulk upload of the zone file data and get the reports. I mean, how many systems like this exist? Because it's like it's been seven years now since we launched this first product and I kind of lost track, just a rough number.

ALAN GREENBERG: If I may intercede, we have about seven speakers in the queue. We have 15 minutes left. I put a two minute timer on, but if everyone uses it, there's no time for answers.

GREG AARON: Okay, just briefly. Registries and registrars use different data sources, and some of them are very well-versed in what's going on in their spaces and some are not.

What I will say is they need to budget for this kind of information and mitigation, and make it part of their budget and the bottom lines. And if they want data, there are various places they can go to get it. What we need to do is make sure everybody does have access.

ALAN GREENBERG: Olivier.

OLIVIER CRÉPIN-LEBLOND: Thanks very much, Alan. And this is one of the topics that, with EURALO, EURALO Chair. This is one of the topics that gets me actually very upset at ICANN and has for many, many different years because this very topic is actually something that deals with harm to end users, and it's something that, unfortunately, I feel ICANN is not doing enough in.

The ALAC has been on record for trying to work out details, for example, in the sensitive strings with the public interest commitments for sensitive strings to, perhaps, have an enhanced amount of security around them, verification of registrants, and so on. And this was blankly pushed out by contracted parties who just want to be able to sell these strings to whoever they want and be in full control of who they want to sell things to.

But in the meantime, I've done a lot of tracking of the Anti-Phishing Working Group, APWG. And I do commend you on that and actually, this very morning, I think, just by coincidence, I sent a circle ID article which points to your latest reports that shows that the situation is only getting worse.

Two years ago, in 2015, when we were in discussions with contracted parties about the use of new gTLDs for phishing and for malware and so on, we were told, "There is no such thing. None of that is happening."

We knew it was going to happen. It's there. What can we do about this?

I've sent something on the new gTLD Working Group of the ALAC and asked whether there should be some kind of provisions made in registrar agreements for making sure that takedowns get done within a certain amount of time, so Service Level Agreements, and I've seen that. And the great majority of them take down websites really quickly, but it looks as though there is a small group where all of those domains are the bad ones and from a specific country that you mentioned in the report. What can we do about this?

GREG AARON: Mitigation, as I said, is something that registrars and registries need to budget for. But sometimes that doesn't happen. And the competition in this space right now is intense, especially if the domains are being sold so cheaply.

What we see is that abuse migrates from place to place over time, and that often, I think, depends upon the price.

Most of these domains that are involved in these crimes don't, aren't a string that is related to the activity. Like most phishing domains don't have the name of the company in them. The phishers use whatever; they don't care.

Again, where I really worry is where large batches of these domains occur and happen over and over again. To an extent, I think this is a compliance problem. We have certain tools. One question is do we need anymore, such as a takedown SLA?

There are some challenges trying to create such a thing because even some of the really responsive, knowledgeable registrars, for example, might have some trouble meeting those. So I think Compliance is one of the best tools we have. I don't know if we're leveraging it as best we could.

OLIVIER CRÉPIN-LEBLOND: Three words. It's a bloody, damn shame.

ALAN GREENBERG: Very well done. Next speaker is Harold. I'm using a one-and-a-half minute timer with speaker, with alarm. Harold.

HAROLD ARCOS: Thank you. I'm going to speak in Spanish. Thank you, Greg, for the presentation.

In this opportunity, I think it would be necessary to remember what happened yesterday and I want to pose a question that many users to me in other ICANN meetings. Yesterday there was a massive attack, and this is the second one after the WannaCry attack. And there were many countries reporting attacks, such as India, Russia, and some companies from the UK.

And in some other occasions, users ask me what ICANN does in these regards, if they are responsible, if ICANN is responsible for enforcing policies in the contracts.

So in your presentation, you are showing that we have policies for that, but there is no precise knowledge about what is happening if this is not met, I mean, how we face in ICANN, these situations. What can you tell us about this? Where can we find information or are we only going to manage the information that is put in the contract? Thank you.

GREG AARON: Hola. The attack he is referring to is a ransomware attack, which started spreading the other day. It's a malware attack. I haven't researched how many domains this is using or exactly how it's managed, but this is typical of these waves of attack.

And again, the mitigation happens two ways. One is the registries may get called because security companies will be calling them. They may also find out about this information if they are monitoring block lists or sometimes called black lists. And these are available. Some are free. Some are commercial.

There are sites that list these domains, and some registrars and registries use those in their monitoring. And they can shut those domains off if they're used for these kinds of things.

Some of this activity is hard to detect. Some of it is spread all over the place. The criminals who are doing this kind of activity will have law enforcement after them because it's widespread and it's international. So it may take some time for law enforcement to find them, but that's the kind of criminal that law enforcement goes after.

HAROLD ARCOS: Apart from law enforcement agencies, where does ICANN find information when these things happen? I mean, the law

enforcement agencies have certain tasks to fulfill. They have some activities within countries, but where does ICANN search for information in order to do the follow-up?

GREG AARON: Sometimes the Security Team at ICANN gets information about this from outside parties, especially if somebody needs assistance finding the right people to talk to. For example, if somebody needs a contact at a registry, the ICANN Security staff will facilitate that contact.

ICANN is also starting to receive some consolidated data about this and I'm helping as a contractor do this.

ALAN GREENBERG: Later on our agenda. Curiously, the next questioner is Dave Piscitello.

DAVE PISCITELLO: Hi, this is Dave Piscitello from ICANN Security Team.

Greg, I was wondering if you could comment on amplification off a domain name into large numbers of URLs in many phishing attacks because the domain name is only one vector, and in many cases, it represents only a fraction of the actual number of attacks.

GREG AARON: Okay. So Dave's talking about some of the domain name, some of the games that cyber criminals can play and one of the things they can do is they can break into someone's web hosting and they can put lots of different things on your domain name. They can put it in a sub-domain. They can put it in a sub-directory.

They can also break up the DNS for your own individual domain names sometimes, and that means they can send people to lots of sub-domains all over the place.

So one domain name can actually be used to perpetrate many different things, or multiple attacks, and that's important to know.

What we also have to realize is that some domain names support multiple services and there are entire companies that are dedicated to selling sub-domain names, so you can go get a sub-domain name. And you don't want to take down that main domain name because then you'll kill all the services that are running on it, many of which may be completely innocent.

So that's one of the challenges in responding to these kinds of problems. You don't want to tell somebody to do something that's going to cause more harm than good.

ALAN GREENBERG: Thank you. We are over time. We have two more people who have asked to be in the queue. We have Kaili and Alberto in that order, if you could be very, very brief because we will have to have a hard stop eventually and we have two more really interesting topics.

KAILI KAN: Thank you, Alan. I will be brief. Kaili Kan speaking. As I'm on the CCTRT, Competition Customer Choice, Consumer Trust, and Consumer Choice Team, Review Team for the new gTLD. Okay, I just wonder if this DNS abuse has… Does the New gTLD Program over the last few years had any effect or impact on DNS abuse, and also, one ongoing discussion within ICANN is the so-called subsequent procedures for new gTLDs and whether there is anything to prevent more DNS abuse could be included in those policymaking. Thank you.

GREG AARON: So one question we can ask is did the new TLDs create more cyber crime or did it enable more cyber crime? Or you could ask were the new TLDs a place where a lot of existing activity migrated to? Those are two different ways to look at it.

KAILI KAN: Actually, that is what I'm asking you.

GREG AARON: Yeah. We see some evidence that the new gTLDs have attracted significant parts of this activity. We see it in phishing and we see it in spam, and we're starting to measure that.

What that means is you have to then work with those operators, registries and registrars, and you have to understand why it went there. I think price, low price is certainly one issue.

The other question is do the new TLDs create more cybercrime? I don't think that's the case. There are plenty of domain names to buy in the existing old gTLDs and ccTLDs. Criminals can get as many domain names as they want and they don't have to get them in one sector or another.

What we have to think about is does the new TLD program, for instance, it has created more operators. Is it possible to get to them? Are they all doing a good job?

ALAN GREENBERG: Thank you. And lastly, Alberto in Spanish.

ALBERTO SOTO: I speak Spanish.

I do understand that you should be working in coordination with the GAC. Okay, so I would like to know what does GAC says in

relation to your notices or in relation about what you say about these attacks because I understand that GAC, not ICANN but GAC, each country can go and take measures, and positive measures. Countries can take positive measures and effective measures.

GREG AARON: Within the GAC, there is something called the Public Safety Working Group, which is composed of law enforcement agents and regulators, among others. And they advise the GAC about what advise the GAC should give to ICANN about cyber crime and domain abuse issues.

And they've been asking ICANN to implement certain policies, like WHOIS accuracy policies, monitoring policies for abuse. And so they're involved in getting the GAC to give advice to ICANN, and then they go back and they advise their individual governments. And this group is important.

ALBERTO SOTO: So my question, I do understand what you say within ICANN. My question is from your group, do the GAC, what do you recommend for countries to do, not ICANN?

GREG AARON: One thing that countries can do that will really help is for them to put in place means for their law enforcement authorities to cooperate and trade information with law enforcement in other countries because, of course, Internet crime goes across all the international boundaries. Law enforcement agents need to be able to talk to their counterparts in other countries and so they need laws that allow them to do that.

If they can't talk to each other and they can't exchange information easily, their hands are tied.

ALAN GREENBERG: Thank you very much. We have to draw this session to an end. I'd like to thank Greg. I think this has been both informative and fun. Thank you very much. And thank you for all the good questions around the room.

GREG AARON: I know we ran out of time. If anyone has questions later, please do find me. I'd be happy to talk with you.

ALAN GREENBERG: We may talk again.

We'll now have Jonathan Zook and members of the CCT Review Team. If other, the CCT Review Team members, there should be

enough seats around the table. Please find one and make yourself at home.

I don't think we need too much introduction of the topic. It's something we've talked about regularly and we have Kaili in the room who has also been giving us regular reports.

Jonathan, do you want to introduce the other people around the room, and then give us a summary of what's going on? We hear you had an interesting weekend.

JONATHAN ZOOK: Yes, thanks, Alan. And thanks, everyone, for taking the time to meet with us and talking about consumer choice, trust, and competition.

I'm Jonathan Zook from the Innovators Network and Chair the CCT. To my left here is Lorraine Kapens from the FTC who is Chair of the sub-team working on trust and safeguards.

And then Jordan Buchanan there from Google is the Chair of the team working on competition and choice. And then I was sort of also the Chair of a sub-team working on the application and evaluation portion of the review as well.

We did have a face-to-face meeting this past weekend and talked about a number of different issues, but the focus was on

two things primarily, three things. One is the public comments that we received, including the ones from ALAC, and how best to address them. And we tried to look at what some of the high level issues were in plenary and then we'll be spending the next couple of weeks in the sub-teams trying to come up with specific responses to the comments on each recommendation.

We also got a presentation on interim results of the DNS abuse report that we commissioned and so we'll be happy to talk about that as well. And oh, Drew Bagley is here who is sort of our chief henchman on DNS abuse. And so, he'll be able to talk a little bit about the results there and we expect a final report in mid-July.

Finally, we talked a bit about a survey that was done by INTA of some of its members about what some of the costs were to trademark holders in the New gTLD Program and how to integrate some of those results into our work.

And another issue from the ALAC comments that is probably worth talking about is parking now that we've had some discussion about it as well over the weekend.

So I don't know what the best way is to proceed. Lorraine, can I hand the talking stick to you to talk a little bit about some of the safeguards and trust things we talked about in lieu of the ALAC comments?

LORRAINE KAPENS: Sure. Oh, sorry.

ALAN GREENBERG: Just to put it in perspective, we have a half-an-hour. In theory, we have 20 minutes left. We can go a little bit over, but not a lot because we have another speaker on domain name abuse again.

JONATHAN ZOOK: All right. [inaudible] topic.

ALAN GREENBERG: Sorry to interrupt.

LORRAINE KAPENS: Always good to get guidance on timeframes and I will keep it pithy. We are still going over the public comments we have received, but we certainly appreciated the comments from the ALAC, which were quite supportive regarding the consumer trust and safeguards recommendations. So we really appreciate that and we share your concerns over the lack of, as useful as it might be in the ideal world, information about the levels of trust that consumers have for the DNS and new gTLDs in particular. We're hoping that we're going to get better information as a result of our recommendations and we're working towards our goal of

getting out the final report so that this process has an end point, hopefully by Abu Dhabi.

But two things that are going to be on your radar screen, I hope, for opportunities for public comments will be the new parts of our report which will be reflecting the results of the INTA study – and perhaps, I shouldn't say "even more important" – I'm particularly focused on DNS abuse, but it may not be more important to you. But I will highlight that the final version of that study is going to be out in July and will incorporate the results from that into our report and those new portions of the report will be put out for public comment. And I'll stop there.

ALAN GREENBERG: I can say with some confidence, it will be of interest to us. Back to you, Jonathan.

JONATHAN ZOOK: Why don't I hand the mic to Drew for a moment, just to talk a little bit about the interim results from the DNS Abuse survey. Just having heard the last part of your last presentation, I think we're going to start to see a lot of studies with similar results vis-à-vis movement of some abuse activity in the new gTLDs rather than an overall growth in those activities.

And in fact, there might even have been some dip in those activities as a promising result of some of the safeguards that were put in place in the new gTLDs. But I'll hand it over to Drew.

DREW BAGLEY: Yeah, it sounds like this is a very appropriate time to speak about this topic here since we're sandwiched in-between two abuse discussions.

HEIDI ULLRICH: I'm terribly sorry to interrupt. Just to remind everyone, we've got English, French, Spanish, and Arabic interpretation, [inaudible] for the transcript, so if you could please say your names every time you speak so that remote participants as well as our interpreters can identify you on the other channels. Thank you.

DREW BAGLEY: This is Drew Bagley from the CCT Review Team, and so right now, we only have a preliminary report and we'll have a lot more information in about a month. But from what we can see so far that is most interesting to us with regard to our mandate is that, in fact, abuse rates as a whole have stayed the same with the introduction of the new gTLDs. So abuse is going up in total

numbers as registrations are going up, but that ratio is actually staying the same.

And as Jonathan was alluding to, what we're looking at is that it appears that there is some substitution going on where people are moving from legacy gTLDs to new gTLDs for certain types of abuse. Despite all of these new safeguards going hand in hand with the new gTLDs from what we saw with the preliminary results, there did not appear to be any overall prevention of abuse in the new gTLDs. But we will have a lot more data and can draw a lot more inferential analysis soon, we hope.

Another interesting thing to note, though, is that the types of abuse that are more prevalent in new gTLDs from the initial report are, first of all, spam. Spam is more prevalent in the new gTLDs than in the legacies. And so it looks like there may have been a migration maybe due to price, we're speculating, but we don't know.

And then another interesting one is that new gTLDs tended to have more maliciously registered domain names instead of merely compromised websites being used for abuse. And so that's another phenomenon that we hope to delve into a bit more next month.

ALAN GREENBERG: Thank you. We'll hold all questions until the speaker is finished. Holly will be first. Jonathan?

JONATHAN ZOOK: Okay, so the only other high-level topic so we can get to Holly as quickly as possible – I know how hard it can be to hold onto a question in the afternoon, the second to last day of a meeting – is parking.

And there were sort of two components to parking. One was its potential contribution, as you mentioned in your comments, or I should say, as we mentioned in our comments to DNS abuse. But then there's also the notion of parking as a counter indicator, if you will, for the competitive effect of the new gTLDs.

And so, I don't know if Jordan, you want to speak a little bit to the competition side of parking and Drew, perhaps, briefly to the DNS abuse side of parking, both of which is lack of information more than a lot of information. But Jordan, you want to say something?

JORDAN BUCHANAN: Sure, and maybe in the interest of time, I'll just briefly address both topics because Drew and I have been chatting about this quite a bit. Jordan Buchanan from Google, again, just for the record.

We did take note. I think, in the ALAC comment, you pointed out that the rate of parking in new gTLDs is quite high. It's in the 60% sort of range. It varies a little bit month to month depending on what source you're using and what's going on in their methodology and so on. But it's definitely a significant number.

So we've spent the past couple of months trying to do two things since the initial report was released, first of all, trying to understand how that compares to parking rates in the legacy gTLDs because it could just be like all TLDs have a lot of parked domains and this is not a new behavior, which turns out to be partially correct.

We actually were able to commission some data from nTLDStats in the same methodology that we were looking at for the new gTLDs and saw that the parking rate for legacy gTLDs was, I want to say 48%, but in the high 40% as well. So it's also very high, but not as high in the new gTLDs.

So then we sort of thought, "Okay, well, what does that, what to make of that differential?" And so we've been looking at it from two perspectives.

Number one is how does it affect our findings around competition? And the other is what's the effect potentially on DNS abuse or consumer trust?

On the competition side, there was a hypothesis within the Review Team that possibly, if there are a lot of parked domains, they wouldn't renew as much as if they were used for things. And so therefore, a TLD with a lot of parked domains might actually, it might look really popular today but if you sort of project forward into the future, that might not be a stable base to sort of ground our findings on competition on.

So we actually went through some ICANN data on renewal rates and tried to find a correlation. Was there any correlation between parking rates and renewal rates? We did a very cursory study of this just to test the hypothesis and in our initial test, we were not able to find a correlation between parking rates and renewal rates.

So, so far, our finding on the competition side is we don't understand if there is any relationship between these parking rates and potential effects on competition. We don't have any other hypotheses to test right now, although that first one about parking versus renewal rates should probably be tested more robustly.

On the consumer trust side, in the DNS abuse study, it's not in the draft report but the authors have indicated to us they've taken a look at this. There is a soft, a very soft correlation, at the lower bounds of what's statistically significant between parking

sites, which is a somewhat more narrow term than what we usually think of as parking, which also includes domains with no name servers and domains other turn errors in addition to the ones that actually have the parking pages with ads and such on them. But those parking pages themselves of TLDs – a lot those – tend to have slightly more abuse associated with them than TLDs with less parking sites.

We don't understand why. We don't understand if it's the parking sites themselves of it's just some other factor of the TLD, like those happen to be TLDs that are really cheap, for example. We don't really know, so we need to do a more deep analysis onto that to see what the actual vectors are, which we expect to be able to do before we get to the final DNS abuse report and our final report.

ALAN GREENBERG: Okay. We'll do a 90-second timer, please, and the first speaker is Holly.

HOLLY RAICHE: What is the kind of abuse that is more common with the new gTLDs, and do you have some kind of idea why that would be possible or why that's happening? Thank you.

DREW BAGLEY: Thank you for the question. From what we've seen so far with the preliminary report, the only one that we know of – well, I guess there's two types, but the only type would be spam. Spam is much more prevalent in the new gTLDs than the legacy gTLDs.

But then the way in which abuse is conducted was the second thing I was referring to and that's the data shows that the new gTLD zones have more domain names that were registered for malicious purposes. Whereas the same type of abuse, whether you're talking about phishing or malware hosting or botnets. Instead in legacy gTLDs more commonly that's from a legitimate website being compromised. Therefore, that could have more to do with the hosting. Or you could be dealing with domain name hijacking and a bunch of other scenarios.

So that's what we've seen with the preliminary data, so we don't know yet why. One theory we're operating with, it could be price. Maybe there are a lot of specials going on in new gTLDs. Like anyone else, bad guys like a good deal. So that could be part of it, we think.

HOLLY RAICHE: Why did you say the registered for? Are you basing that on evidence?

DREW BAGLEY: It's based off of a model that the researchers came up with – the researchers who did the DNS abuse report. They looked at how many days a domain took to go bad, essentially, after it was registered and factored in a few other variables with that. They explain it in the preliminary report. If you read that, that'll offer a better explanation than me off the top of my head. But that's what they're looking at to come up with this, and that model was based off of a study from a couple years ago that found this commonality between domain names that were registered by people with bad intent to DNS abuse.

ALAN GREENBERG: Thank you. I'm next in the queue. Question for Jordyn. If I register Hilton.hotels and redirect it to Hilton.com, do you consider that parking? And if so, have you attempted to measure what percentage of the parked domains are of that sort, that ilk?

JORDYN BUCHANAN: Thanks. We have a very expansive definition of parking that roughly is: anything that does have content directly hosted on that domain. It includes no DNS, DNS errors, http errors, parking pages, and redirects. So to answer your question, yes, that would count as parking using our current methodology. It is a relatively small percent of total "parking" that's attributable to redirects. I think it's about 4%. Maybe it's 4% of the total, so it'd

be about 8% of the parking or something like that. But it's not the majority use case.

ALAN GREENBERG: Thank you. Next I have Satish.

SATISH BABU: Thank you, Alan. I'd like to know if there is any evidence of IDNs (Internationalized Domain Names) being especially susceptible to malicious use.

DREW BAGLEY: For the preliminary report, we don't have any data on that, I don't believe. That's definitely a question that we're very interested in, of course, because of, obviously – I can't think of the name – these homonym-style attacks, where you're using different characters to look like another character set. So that's something I don't know if we actually have the data for. I don't know if the researchers were able to look at those trends or not, but I can certainly get back to you on that.

ALAN GREENBERG: Thank you very much. Olivier?

OLIVIER CREPIN-LEBLOND: Thanks very much, Alan. Forgive me if I might sound a little provocative here. We've had the visit just before you of David Aaron's from the Anti-Phishing Working Group. They're collecting a lot of data on all these issues, and I see, of course, now with this whole collecting of more data and things – what is ICANN doing about this? We're going to have all this data. Is this data going to be actually used for doing anything? Because at present, it really reads like a philosophical thesis about the sexual life of the ping pong ball. What I mean is that there's lots of data, but for what?

ALAN GREENBERG: You may want to reserve that question for someone from ICANN, but I'll let Drew answer if he'd like.

DREW BAGLEY: I don't shy away from questions. Not shying away from questions. I think what you're saying is something that we're all very concerned about in general, and that's why a lot of what we're doing with our recommendations in general is calling for a data-driven approach to policy making at ICANN. Specific to the DNS Abuse study, this is the first time that a comprehensive analysis has ever been undertaken. APWG's data feeds were used as part of this analysis. APWG does great work and they do quarterly reports, but this was the first time that there had been

comprehensive, historical analysis that looked at every single zone.

Our intention with that was so that we could measure whether or not the safeguards put in place to mitigate abuse were effective or not and whether or not consumer trust and perceptions about new gTLDs were actually in line with the reality of whether or not they could trust them, not trust them, or trust them just the same as any other domain name.

Once we have this final report, we're actually going to hopefully be able to make some policy recommendations based on the data. So for us, it absolutely not a philosophical thesis on the sex life of the ping pong ball, but instead it was more of we're putting together a strategy on how to better play ping pong.

LAUREEN KAPIN: Just to add on, the ideal scenario here is that, once you have specific data that correlates some type of behavior to some type of ill, whether that is phishing or some other sort of DNS abuse that that will feed into the policymaking process. For example, the GAC could take a look at that data and say, "Wow. We really need to consider the relationship between price and DNS abuse," or, "Perhaps we need stronger contract provisions that have quicker and more effective monitoring of DNS abuse and mandates certain responses." Those are concrete ways I can see

this data being used, so it is not just an exercise in data collection, which is all very fine and well. But it's where the rubber meets the road in terms of behavior and how that's monitored and how that's sanctioned that really matters here.

JONATHAN ZUCK: Sorry. As I bounce this idea around in my head a little bit more, the ping pong ball – there's a number of different issues at play here. One, as Laureen mentioned, is price. There's an expression in English – the third rail – which is about the subway system rail that's electric that should not touch. Price is sometimes that inside the ICANN community. So I think it's going to require a lot of discussion because ICANN justifiably stays away from seeing itself as a price regulator.

We also see that the highest propensity, as Drew mentioned, is in the area of spam. Spam is not uniformly considered illegal around the world in different jurisdictions, so it's not one of the things called out specifically as DNS abuse in the contracts today. Even though we view it as the gateway drug to phishing, spam in and of itself is a little bit more difficult to assess.

I will say – Drew, do correct me if I'm wrong about this – but, if you look at the charts in the DNS abuse report, they do show not a significant but a small dip in DNS abuse associated with the New gTLD Program that may reflect some success of the

safeguards that were put in place because the list of safeguards is long in the new contracts. If we start to see that some specific practices of some specific registry operators are leading to lower amounts of DNS abuse, then that'll give us something to point to empirically as practices that should be adopted across the board going forward.

ALAN GREENBERG: Thank you. Our last speaker is Sebastien Bachollet.

SEBASTIEN BACHOLLET: To be short, I will do it in English. Did you have done any comparison with the previous round about the data you gathered?

JONATHAN ZUCK: In what aspect?

SEBASTIEN BACHOLLET: Whatever aspect. Did you have any data from the previous round? Because I guess your answer will be "We don't have any data from the previous round." If you have, then which comparison you have made with the previous rounds?

JONATHAN ZUCK: Jordyn, go ahead.

JORDYN BUCHANAN: Sebastien, thanks. It depends on the topic, I think, and roughly whether there was a previous effort to gather data on the previous rounds. In general, as I think you've heard from us, there hasn't been robust data collection by ICANN prior to the preparation for this review. At the time the 2000 and 2006 rounds were occurring, there wasn't a huge amount of data collection by ICANN in a lot of these metrics.

However, in some cases there were third-party studies – academic studies and so on – in particular areas. In many cases, we have looked to those previous studies to try to understand the dynamics. Most of them oriented around the competition side of the equation. So we could look to see for example – one of the things that we looked at is how quickly the TLDs grow – for example, .biz and .info – in the 2000 round and how that compares to the new gTLDs.

I don't think we've seen any significant differences in dynamics, at least on the competition side, from the previous rounds to this round. The big change is that this time we looked at the set of TLDs as a whole. In the past, we were looking to say, "How does .biz compare to .com?" or something like that, and our analysis this time really centered on "How does the collection of

all these TLDs together compare to the legacy gTLDs?" because that represents, in our mind, the consumer choice. We have all these new choices versus one particular TLD trying to establish itself competitively. But we did look at some of the data about the previous expansions.

SEBASTIEN BACHOLLET: Just one addition, and it will be the last time I will say that. There were data collection for the 2000 round. I can never publish them. When you talk about parking, for example, it was done. It was done and the report was sent to ICANN on that but never published. I think it could be interesting to try to find them. If I am the only one to have this, maybe somebody from ICANN can ask me about it. Thank you.

ALAN GREENBERG: Thank you, Sebastien. Thank you, Jonathan and everyone. We're doing really well this afternoon so far, at least in my estimate. We've had two for two and a really great session. Thank you very much.

JONATHAN ZUCK: Great. Thank you.

[DREW BAGLEY]: Thanks for having us.

ALAN GREENBERG: Next up we have Dave Piscitello will talk about - if you look in your agenda, you'll see that he's talking about the TTFKAD. That stands for The Tool Formerly Known As DART. Dave will give you a bit of history on that.

You're not allowed to click yourself.

DAVE PISCITELLO: Oh my Lord. Okay.

ALAN GREENBERG: Our union agreements require that other people do the clicking.

DAVE PISCITELLO: Okay. That's fine.

ALAN GREENBERG: We also don't have a laser pointer. I've never understood why.

DAVE PISCITELLO: Let me introduce myself for the record. I'm Dave Piscitello. I'm the Vice-President of Security and ICT Coordination at ICANN. While we're waiting for the slides, I will explain The Tool

Formerly Known As DART, which is now called the Domain Abuse Activity Reporting System or Reporting Platform.

Just about 30 minutes before I was boarding a plane to come to Johannesburg, I received an e-mail from our Legal Department saying that ICANN had received a cease-and-desist letter for blocking the use of DART as an acronym because another organization claimed copyright over it. So we have changed the name to – next slide – oh, I can next-slide it now –

UNIDENTIFIED FEMALE: No, actually, Dave, sorry. It doesn't work because Yesim has to do it.

DAVE PISCITELLO: Okay. So the project now moving forward is the Domain Abuse Activity Reporting Project. We'll call the system the DAAR system. It turns out that, after my search, I found that "daar" meant "there" in South Africana, and it means "house" in Arabic. Our Legal Team did a search for us and there seems to be no other conflicts, so I'm hoping that next time I see you the name will remain the same. Next slide, please.

Let me explain what we've done. We are building a platform to report domain name registration and abuse across TLD registries and registrars. I in particular have been doing this for

many more years than I've been at ICANN – and I've been at ICANN for twelve years – listening to very large outcries of indignation similar to those we've heard here about how much abuse there is.

Very little reliable data have been presented. Much of the data either comes from a commercial source with limited resource or comes from academia with, again limited resource.

One of the things that we decided to do was try to be more comprehensive and to be as scientific as possible. We study all the TLD registries and registrars from which we can collect zone and registration data. Currently that means that you're studying around 1,241 TLDs.

We also looked in the literature and looked at many of the commercial reports and noticed that, in most cases, you only saw a limited number of feeds or reputation providers' data were incorporated into that study. We have a very large number of reputation feeds. I'll identify them for you later on in the presentation.

Much like the CCT effort, we want to have historical studies. We will be able to provide day and time, but our data will take us back for all these TLDs to January 1st, 2017. So we're moving forward. We now have a very, very large database for historical purposes. Obviously, collecting data from the TLDs and

registrars was actually going to be easier than trying to get reputation data several years past because many of the providers don't keep it. So we are actually going to be a rather unique repository of that reputation data.

We're also studying multiple threats, and we're extensible. We began with the three threats that the GAC identified in their original communiqué, which were phishing, botnet, and malware. We don't do pharming because there are very few reputation or processes that allow classification of pharming. It's actually a form of attack, not a security threat.

Since we believe that spam is a very important part of this picture and the GAC in Hyderabad actually had a correspondence with the Board, saying that their previous security threats were examples and that they also were interested in spam, we expanded our study to include spam as well.

What we intend to do is create a set of data that is unbiased. We're gathering the data that can be obtained by any of you, and we're using the same means that anyone else could gather it with. We'll be publishing a paper about our methodology shortly. What we intend to try to do is emulate good scientific practices. We expect that, if you were to go and do what we do

with the methodology and the data that we use, your results would be similar to what we found. Next slide, please.

You may have heard of the Open Data Initiative. This is another project, an activity, in the Office of the Chief Technology Officer. We are attempting to facilitate access to data that ICANN organization or the community creates or curates. In anticipation of the question – what are we going to do with the data? – we expect to be able to publish the data.

The project uses only public, open, and commercial-sourced data. We get DNS zone data. We use WHOIS data. We use open source reputation data. Then there were certain commercial feeds that require us to pay a monthly license that are part of our feeds.

The one constraint and limitation that we know we have to deal with in the future is that, where we have data that we're able to use in a derivative manner but not in a proxying manner – in other words, where we cannot share the data because we have a direct license – we won't be able to publish that data unless we have a different kind of licensing arrangement. But if that's what the community wants, then we have to look into that. Next slide.

What's our goal? This is not a name-and-shame project. This is not a "point at people and say, "You're bad" project, partly because, as Greg began discussing and I'm happy to talk about

in my time slot, identifying who's at fault, whether it's the hosting company, the registrar, or the registry is not as simple as many people tend to imagine. Trying to find the right action to take that falls into compliance or whether it's an action that has to be taken in a single or multiple jurisdiction by law enforcement is something that I work on almost every day as part of my day job. I can tell you that it is a very, very non-trivial, complex process.

What we hope to do with our data is to give the community data that will support the Policy Development Process. I think one of the reasons why policy lags in certain places is that there aren't enough data for people to make informed decisions. What we hope to do with the DAAR project and moving forward in more big data projects like these is to give the community ample data to make an informed decision about a policy or about the unintended consequences of a policy and ways to resolve that. So that's the goal here.

What can we do with the data? We hope that we'll be able to identify threats reported at a TLD registrar level for all TLDs for which we can obtain data. It would be nice if one of the things that we can do in the future is provide threat lists. That's something on our scope.

We want to be able to track security threats because point-in-time is so fluctuating in this space. You can take one of the TLDs, for example, 60 days ago, they had 160,000 registrations. This month they have 40. So there's a large variation in some of the ways that registration behavior is proceeding.

The same thing is true for spam. There are spikes in spam that are associated with massive campaigns. There are drops in spam associated with massive takedowns. So a single day doesn't help you.

However, a histogram that shows you 30 or 60 or 90 or 365 days actually starts to give you some ideas of what's happening. We'll have data for each of those days to try to go back and try to understand or discuss with an operator and say, "What happened on this day?"

We think that this is going to create an opportunity for ICANN staff to work with the ICANN community, particularly the contract parties, to consider different ways of managing their reputations, managing their anti-abuse programs, and reconsidering, if necessary, their terms of service.

Obviously, one of the most important features is to study malicious registrations behaviors because this is a very large part of the criminal ecosystem. The ability to register domains and then use them for various malicious or criminal purposes is

a very, very big problem for us. As Greg mentioned, it appears, at least in the last year, that the number of domains registered maliciously, at least for phishing, has gone up. In fact, I think their study says it's tripled. So that's a concern and we want to understand why that's happening. Next slide, please.

So let me explain how we use domain and zone data. We collect the zone files from the registries using the Centralized Zone Data Service or legacy zone transfer. We have approximately 195 million domains that we have in our corpus at the moment and 1,241 top-level domains.

We have actually been talking to a number of the ccTLDs – well, I made this presentation at the DNS Symposium in Madrid, and six TLDs came up to us and asked if they could participate. So we're discussing with them. While I was here this week at ccNSO, I had three more. It would be really nice if they all participate. Then we would have a very accurate picture of the entire name space. Next slide, please.

We use WHOIS. Fortunately for us, we get to dodge the entire GDPR issue, hopefully, because the only thing that we use from WHOIS is the sponsoring registrar. This is the way we associate the domain name with the registrar portfolio.

Okay. All the other WHOIS data, point of contact data, is not important to us. We might find a use for creation date and

expiration date to do some other analysis of the standard amount of time or the median amount of time that a name that is malicious remains alive. We may look into things like when was the domain was first observed in the DNS, but we don't need anything related to personal identifying information.

We don't use every single name in a registry's portfolio. We only use the names that can resolve. Our philosophy is that a security threat can't be perpetrated against a user or executed if the name can't resolve to an IP address. Next slide.

I mentioned that we use many threat data sets. We have spent probably the largest amount of time in this project trying to decide what threat data sets we would consider and incorporate. We have 20 data sets from 12 reputation providers. We've gone through I can't tell you how many academic articles, commercial reports, looking at the vendors, talking with the vendors that sell the data to try to understand their methodologies and their practices.

I feel that Greg Aaron and I who have put together this lists have chosen lists that have a high history of accuracy. Collectively they give us global coverage and false positive rates.

One of the things that we want to do with this project that differs from a lot of other projects is that we're trying to bring the lens of what the user community sees or the enterprise sees the

name space to be, through their security systems, through their defensive measures that they take to prevent spam, phishing, malware, and the like.

So we're not trying to come up with new data. We're trying to take the data in a composite and trying to allow our community to see how other communities external to ICANN see the ecosystem.

Okay. We've built this project to be extensible, and we expect that they we'll be able to add new data that we feel is reliable. If, for example, we find at some future time that one of the lists that we're using now becomes less reliable and less accurate, we could conceivably drop that. There are circumstances where historically a list has begun as a research project. It's had some steam and momentum while it had some funding, and then the project gets neglected. So the reputation of the list itself diminishes over time. If that happens to one of our lists, we will dump it. Next slide, please.

I want to make certain that it's very clear that ICANN is not now in the business of creating block lists. We're not in the business investigating phishing or investigating malware or investigating spam. We are using what reputation providers feed to us.

We use domain and URL abuse data, and we curate that. We also rely on the reputation providers that we believe do the best

curation of their data. We manage a number of counters. We count security threat domains and we classify them into the four categories that I mentioned before: spam, phishing, malware, hosting, and botnet. We also count total abuse domains, and we have a running total from day one. So that will be a 365-day running window as we proceed and get into the point where we have years of data.

We can also automatically generate histograms. We generate charts that compare the various TLDs – the new TLDs against legacy TLDs against IDN gTLDs and hopefully ccTLDs.

When we count a domain, we deduplicate. Obviously, when we're using multiple lists, there are occasions when we actually would find a domain listed twice. We only count it once. So we feel like we have a fairly large, accurate set of abuse domains. Next slide, please.

This is our current list of reputation data sets. I apologize, but SURBL belongs on that list and it doesn't appear. We've chosen these lists to assure that we have at least two lists for each of the security sub-classification or threat sub-classifications that we're looking for. We also chose them because they support the classification mechanisms that we're seeking. Next, please.

We've received a lot of curiosity and in some cases worry that we're using a lot of data sets and you know isn't that going to

cause a lot of duplication. It turns out that, doing a little bit more research and speaking with some people at Carnegie Mellon who had done some research, their paper – Blacklist Ecosystem Analysis – which was actually done over a period of three years, actually finds that there's little overlap in blacklists.

Part of the reason is that most of the services don't use the same methodologies. Most of them don't spam trap or have collection networks that are completely self-similar but operate in different geographic areas or operate with different partnerships with different Internet service providers. So we felt confident that we could use multiple lists.

We actually began with a list of 86 lists. Both Greg and I did a number of tests, running scripts against a subset of the TLDs. Our results, looking at those 86 lists, confirmed what the Metcalf and Spring paper showed.

So our feeds, I think, represent probably the best in the industry for clarity of process and accuracy. They use the threat classifications that we want. When we started to look at security system adoption among the commercial vendors, there was almost a consensus adoption for many of the ones that we use. We also tried to assess the quality based on the frequency citation in academic literature. Next, please.

I'm just going to go through this quickly and answer no because I'd like to save time for talking. We don't capture all the abuse. We capture a whole lot of it, certainly enough to make accurate assessments. Next slide, please.

In addition to the counters that we have, one of the things that we are experimenting with is some way to measure large TLDs against small TLDs with some way to normalize. We have a simple percentage of abuse as our first metric of sorts. We're going to solicit community input for other metrics that you might want us to pursue, and that's part of the whole process of talking with you and the other SOs and ACs this week. Next, please.

Our percent of abuse is relatively simply to calculate. It's just a fraction. It's the number of domains that were on an abuse list in a TLD on a given day over the number of domains that resolved in that zone on that day times 100. Next.

The registrar is the same. The numerator is the number of domains sponsored by the registrar. Next, please.

Let me show you some of the samples of the kinds of visualizations of our data that we intend to share. This visualization uses that percent of abuse to illustrate that phishing and spam seem to be migrating to the new TLDs or being distributed across all the TLDs rather uniformly, whereas

malware hosting and botnet command and control remain largely in the legacy TLDs. This is a single day and time. Next.

You're missing a slide – there we go. Thank you.

This is some of the more interesting data that we're drilling down and exploring. This is a scatter plot. Each of the blue diamonds represents a top-level domain. Both axis are logarithmic scales because we had a lot of congestion in representation. The line across the scale there represents 0.6, which is the median abuse score. As you can see, a great number of the TLDs, legacy and new, are doing quite well. In fact, what is not depicted on this slide is that, as of May 31st, only 356 of the 1,241 TLDs had one or more incidents. So there are a number of TLDs that have no incidents at all.

The place where we get concerned is above that line – the X axis – that says "10." There about 25 TLDs there that have extremely large percents of abuse, which means that they have the largest numbers of malicious registrations or abuse domains reported in the ecosystem.

One of the interpretations of a slide like this is: why is that the case? And what do we need to do as a community to drive all the abuse scores down closer to the mean or closer to zero? Next slide, please.

UNIDENTIFIED MALE: Are you almost done?

DAVE PISCITELLO: Yes, I'm almost done. This is the last slide. What is the status of the project? We're currently in beta. We had a lot of difficulty collecting WHOIS, and we finally managed to sort out the collection mechanism so that we can keep pace with the number of WHOIS records that we have to process on a daily basis. We are entertaining interests for cc operators to join us, and the biggest part of the talk here, I hope, is to find out from the community how you want us to report this information, what kind of representations or findings or data you want us to present, to whom are we going to report this, and what kind of access to the data we should consider?

I'll stop talking and I'll start answering questions. Thank you very much.

ALAN GREENBERG: I'm going to have to leave in a moment. I'll turn the floor over to Olivier if he'll take it. Thank you. We have three people in the queue. I'm there and I'll ask a very quick question, and then Alberto and Garth, and I'll ask Olivier to manage the queue after that. We are officially out of time, but we're on break, so enjoy.

Can you go to the previous slide?

DAVE PISCITELLO: Well, I can't but –

ALAN GREENBERG: No, I – thank you. You said you're not in the business of name and shame, but I notice there's one registry that's almost 100%. Are we going to have ready access to the data which will allow us to identify who that is? You may not want to name and shame, but some of us might.

DAVE PISCITELLO: This is a question the community has to answer. I am a database. I have a database for you. What ICANN community must decide is how you want us to represent this data and in what forms. Then we have to sit and understand how we build that because we don't have a user interface that anyone can walk up to and type in and get whatever data they want. So we have to generate a reporting system. In order to do that, we need to get an idea of what kind of reports you want us to do. That's part of the reason why we're in beta.

So I encourage ALAC to do the same as what the PSWG did when we met with them. They had a number of ideas of how they

wanted to have the information reported. They're going to pass it up to the GAC. The GAC is going to present it to the Board. I encourage you to do the same thing. If you want to name and shame, then you have to go and justify it. If everyone says that's what they want, then I give you the data and you generate those lists. If you want to name and fame – because I will tell you right now, some of the players that get most criticized by anecdotal information are actually some of the best operators in the field. So I would be more than delighted to actually create a list of the best 25 registrars or the best 25 TLDs as well.

ALAN GREENBERG: Thank you. We have – sorry. We have Alberto, Garth, and Holly. I have to leave, and I'm told the interpreters have five minutes. We continue past that, but only in English. Thank you.

OLIVIER CREPIN-LEBLOND: Next is Alberto Soto.

ALBERTO SOTO: Thanks. I will speak in Spanish. I am aware that this is an extremely complex topic as we have seen in previous presentations, as complex as looking after our grandmother so that she might not die. But it happens that many constituencies are requesting ICANN to look after grandmother. But that same

constituency has to check the temperature. When grandmother dies, the blames goes to ICANN.

So my specific question is, what has GAC said it was going to do with respect to the first objective, which is the identification? Because ICANN cannot work on its own, and the GAC should have some commitment to contribute.

DAVE PISCITELLO: Thank you. I'm not going to speak for the GAC. What I reported to you is that I presented this information to the Public Safety Working Group, and they were very enthusiastic about the opportunities that this data set represented. So they are going to review together, come up with a list of recommendations for the GAC, and then the GAC will then go to the Board with whatever derivative of that list they feel is appropriate for Board action and for community consideration.

So I don't know what the GAC will do because, just like you, I presented to the PSWG earlier this week. I'm optimistic that, with all the enthusiasm, we will be able to do some meaningful reporting.

OLIVIER CREPIN-LEBLOND: Thanks very much, Dave. I have a list, thanks to my predecessor, of Garth, Holly, Harold Arcos, and Ricardo Holmquist, and we're

already beyond time. Is anyone ready to – so no, Holly? Garth, you want to remain on record? Okay. And Harold, do you want to remain on record, or can you speak with Dave Piscitello afterwards?

DAVE PISCITELLO: I'm happy to take this out in the hall and spend as much time as you need to –

OLIVIER CREPIN-LEBLOND: And Ricardo as well – could you speak with Dave directly afterwards?

RICARDO HOLMQUIST: Yeah. No problem.

OLIVIER CREPIN-LEBLOND: It's just for interpreters that haven't had a break. Let's then have Garth Bruen, please.

GARTH BRUEN: Garth Bruen, ALAC North America. I want to echo what Olivier said earlier in that there's been a lot of noise and very little action. Obviously something has to be studied before it can be addressed. I have absolutely no doubt, Dave, that your results

are going to be specific and accurate. I have no doubt in that whatsoever. A lot of faith in you.

The issue is in the follow up. What's going to happen next? Let's assume for a minute that there's an effective enforcement process in the background. Based on what Greg said and what you've said and what other people have said, some of the things that might change are that prices may go up, a large number of domains might be removed from the DNS that are abusive, and mechanisms have to be created in the background to do effective enforcement.

All of these things are going to reduce ICANN's net income. What is ICANN's incentive for doing any of that?

DAVE PISCITELLO: I've been in the domain world for 15 years. I was one of the first 6,000 people getting a com domain. I paid $40 or $50 for it. Before we had a race to the bottom in pricing – I think that, if the industry had stayed at $40 0r $50 a domain for everyone, parties would be making approximately the same money. It just wouldn't be as many domains.

I don't want to get into pricing because I'm just starting to look at that. I hope to have a correlation between some of these numbers and pricing by Abu Dhabi. It's one of my next goals. But

I think it's premature to make assumption that ICANN's revenue or the community's revenue is going to crater because we're going to get rid of abuse.

Let's be honest. We're not going to get rid of abuse. Abuse has been here for as long as a domain name was more attractive than an IP address. If we go and we scatter the roaches, so to speak, we may end up with approximately the same amount of abuse. Everyone's score might be even and everyone might have approximately the same amount of abuse.

As Greg said, we don't really see that there's a significant uptick. We see that there is shifting. What my data tell me is that, not only is there a lot of shifting, but there's a lot of tasting and flocking and migration from one TLD to another. Just anecdotally, I can tell you that the one close to 100 is now down around 50 because they're dumping domains.

So who knows why? There's so much additional information and intelligence that we can apply in the future once we have this data. In a year, we may actually know why. We may be able to have a better insight into countermeasures that we can put in place to prevent these sorts of things. But for the first time, in being here for 12 years, I feel confident that we have some data to start moving forward and doing some of the really meaningful consideration of how to manage abuse.

OLIVIER CREPIN-LEBLOND: Well, thank you very much, Dave, for joining us. Indeed, we've had an hour-and-a-half of horrendous – well, great programs for data collection. But it looks as though the data that's coming in doesn't look that great for things.

UNIDENTIFIED MALE: [inaudible]

OLIVIER CREPIN-LEBLOND: Yeah. For end users. Anyway, thanks to our interpreters for having extended their time, and thanks to everyone here. Now there are main sessions going on, so you're all invited to go and take part of these sessions that take place in other rooms.

Thanks very much, and have a very good afternoon. This session is adjourned.

[END OF TRANSCRIPTION]